

Uniform Stability of a Particle Approximation of the Optimal Filter Derivative*

Pierre Del Moral[†], Arnaud Doucet[‡], Sumeetpal S. Singh[§]

June 14, 2011

Abstract

Sequential Monte Carlo methods, also known as particle methods, are a widely used set of computational tools for inference in non-linear non-Gaussian state-space models. In many applications it may be necessary to compute the sensitivity, or derivative, of the optimal filter with respect to the static parameters of the state-space model; for instance, in order to obtain maximum likelihood model parameters of interest, or to compute the optimal controller in an optimal control problem. In Poyiadjis et al. [2011] an original particle algorithm to compute the filter derivative was proposed and it was shown using numerical examples that the particle estimate was numerically stable in the sense that it did not deteriorate over time. In this paper we substantiate this claim with a detailed theoretical study. \mathbb{L}_p bounds and a central limit theorem for this particle approximation of the filter derivative are presented. It is further shown that under mixing conditions these \mathbb{L}_p bounds and the asymptotic variance characterized by the central limit theorem are uniformly bounded with respect to the time index. We demonstrate the performance predicted by theory with several numerical examples. We also use the particle approximation of the filter derivative to perform online maximum likelihood parameter estimation for a stochastic volatility model.

Some key words: Hidden Markov Models, State-Space Models, Sequential Monte Carlo, Smoothing, Filter derivative, Recursive Maximum Likelihood.

1 Introduction

State-space models are a very popular class of non-linear and non-Gaussian time series models in statistics, econometrics and information engineering; see for example Cappé et al. [2005], Doucet et al. [2001], Durbin and Koopman [2001]. A state-space model is comprised of a pair of discrete-time stochastic processes, $\{X_n\}_{n \geq 0}$ and $\{Y_n\}_{n \geq 0}$, where the former is an \mathcal{X} -valued unobserved process and the latter is a \mathcal{Y} -valued process which is observed. The hidden process $\{X_n\}_{n \geq 0}$ is a Markov process with initial law $dx\pi_\theta(x)$ and time homogeneous transition law $dx'f_\theta(x'|x)$, i.e.

$$X_0 \sim dx_0\pi_\theta(x_0) \text{ and } X_n|(X_{n-1} = x_{n-1}) \sim dx_nf_\theta(x_n|x_{n-1}), \quad n \geq 1. \quad (1.1)$$

*First version: January 2011. Cambridge University Engineering Department Technical report number CUED/F-INFENG/TR.668

[†]Centre INRIA Bordeaux et Sud-Ouest & Institut de Mathématiques de Bordeaux, Université de Bordeaux I, 351 cours de la Libération 33405 Talence cedex, France (Pierre.Del-Moral@inria.fr)

[‡]Department of Statistics, University of British Columbia, V6T 1Z4 Vancouver, BC, Canada (arnaud@stat.ubc.ca)

[§]Department of Engineering, University of Cambridge, Trumpington Street, CB2 1PZ, United Kingdom (sss40@cam.ac.uk)

It is assumed that the observations $\{Y_n\}_{n \geq 0}$ conditioned upon $\{X_n\}_{n \geq 0}$ are statistically independent and have marginal laws

$$Y_n | \left(\{X_k\}_{k \geq 0} = \{x_k\}_{k \geq 0} \right) \sim dy_n g_\theta(y_n | x_n). \quad (1.2)$$

Here $\pi_\theta(x)$, $f_\theta(x | x')$ and $g_\theta(y | x)$ are densities with respect to (w.r.t.) suitable dominating measures denoted generically as dx and dy . For example, if $\mathcal{X} \subseteq \mathbb{R}^p$ and $\mathcal{Y} \subseteq \mathbb{R}^q$ then the dominating measures could be the Lebesgue measures. The variable θ in the densities are the particular parameters of the model. The set of possible values for θ , denoted Θ , is assumed to be an open subset of \mathbb{R}^d . The model (1.1)-(1.2) is also often referred to as a hidden Markov model in the literature Cappé et al. [2005].

For a sequence $\{z_n\}_{n \geq 0}$ and integers i, j , let $z_{i:j}$ denote the set $\{z_i, z_{i+1}, \dots, z_j\}$, which is empty if $j < i$. Equations (1.1) and (1.2) define the law of $(X_{0:n}, Y_{0:n-1})$ which is given by the measure

$$dx_0 \pi_\theta(x_0) \prod_{k=1}^n dx_k f_\theta(x_k | x_{k-1}) \prod_{k=0}^{n-1} dy_k g_\theta(y_k | x_k), \quad (1.3)$$

from which the probability density of the observed process, or likelihood, is obtained

$$p_\theta(y_{0:n-1}) = \int dx_0 \pi_\theta(x_0) \prod_{k=1}^n dx_k f_\theta(x_k | x_{k-1}) \prod_{k=0}^{n-1} g_\theta(y_k | x_k). \quad (1.4)$$

For a realization of observations $Y_{0:n-1} = y_{0:n-1}$, let $\mathbb{Q}_{\theta,n}$ denote the law of $X_{0:n}$ conditioned on this sequence of observed variables, i.e.

$$\mathbb{Q}_{\theta,n}(dx_{0:n}) = \frac{1}{p_\theta(y_{0:n-1})} \left(dx_0 \pi_\theta(x_0) g_\theta(y_0 | x_0) \prod_{k=1}^{n-1} dx_k f_\theta(x_k | x_{k-1}) g_\theta(y_k | x_k) \right) dx_n f_\theta(x_n | x_{n-1})$$

Let $\eta_{\theta,n}$ denote the time n marginal of $\mathbb{Q}_{\theta,n}$. This marginal, which we call the filter, may be computed recursively using Bayes' formula:

$$\eta_{\theta,n+1}(dx_{n+1}) = \mathbb{Q}_{\theta,n+1}(dx_{n+1}) = \frac{dx_{n+1} \int \eta_{\theta,n}(dx_n) g_\theta(y_n | x_n) f_\theta(x_{n+1} | x_n)}{\int \eta_{\theta,n}(dx'_n) g_\theta(y_n | x'_n)}, \quad n \geq 0$$

and $\eta_{\theta,0} = \pi_\theta$ by convention. Except for simple models such the linear Gaussian state-space model or when \mathcal{X} is a finite set, it is impossible to compute $p_\theta(y_{0:n})$, $\mathbb{Q}_{\theta,n}$ or $\eta_{\theta,n}$ exactly. Particle methods have been applied extensively to approximate these quantities for general state-space models of the form (1.1)–(1.2); see Cappé et al. [2005], Doucet et al. [2001].

The particle approximation of $\mathbb{Q}_{\theta,n}$ is the empirical measure corresponding to a set of $N \geq 1$ random samples termed particles, that is

$$\mathbb{Q}_{\theta,n}^{P,N}(dx_{0:n}) = \frac{1}{N} \sum_{i=1}^N \delta_{X_{0:n}^{(i)}}(dx_{0:n}) \quad (1.5)$$

where $\delta_z(dz)$ denotes the Dirac delta mass located at z . This approximation is referred to as the *path space* approximation Del Moral [2004] and it is denoted by the superscript 'p'. The particle approximation of $\eta_{\theta,n}$ is obtained from $\mathbb{Q}_{\theta,n}^{P,N}$ by marginalization

$$\eta_{\theta,n}^N(dx_n) = \frac{1}{N} \sum_{i=1}^N \delta_{X_n^{(i)}}(dx_n).$$

These particles are propagated in time using importance sampling and resampling steps; see Doucet et al. [2001] and Cappé et al. [2005] for a review of the literature. Specifically, $\mathbb{Q}_{\theta,n+1}^{\text{p},N}$ is the empirical measure constructed from N independent samples from

$$\frac{\mathbb{Q}_{\theta,n}^{\text{p},N}(dx_{0:n}) dx_{n+1} f_{\theta}(x_{n+1}|x_n) g_{\theta}(y_n|x_n)}{\int \mathbb{Q}_{\theta,n}^{\text{p},N}(dx_{0:n}) g_{\theta}(y_n|x_n)}. \quad (1.6)$$

It is a well known fact that the particle approximation of $\mathbb{Q}_{\theta,n}$ becomes progressively impoverished as n increases because of the successive resampling steps [Del Moral and Doucet, 2003, Olsson et al., 2008]. That is, the number of distinct particles representing the marginal $\mathbb{Q}_{\theta,n}^{\text{p},N}(dx_{0:k})$ for any fixed $k < n$ diminishes as n increases until it collapses to a single particle – this is known as the *particle path degeneracy* problem.

The focus of this paper is on the convergence properties of particle methods which have been recently proposed to approximate the derivative of the measures $\{\eta_{\theta,n}(dx_n)\}_{n \geq 0}$ w.r.t. $\theta = [\theta_1, \dots, \theta_d]^T \in \mathbb{R}^d$:

$$\zeta_{\theta,n} = \nabla \eta_{\theta,n} = \left[\frac{\partial \eta_{\theta,n}}{\partial \theta_1}, \dots, \frac{\partial \eta_{\theta,n}}{\partial \theta_d} \right]^T.$$

(See Section 2 for a definition.) References Cérou et al. [2001] and Doucet and Tadić [2003] present particle methods which have a computational complexity that scales linearly with the number N of particles. It was shown in Poyiadjis et al. [2011] (see also Poyiadjis et al. [2009] for a more detailed numerical study) that the performance of these $\mathcal{O}(N)$ methods, which inherently rely on the particle approximations of $\{\mathbb{Q}_{\theta,n}\}_{n \geq 0}$ constructed as in (1.6) above, degraded over time and it was conjectured that this may be attributed to the particle path degeneracy problem. In contrast, the alternative method of Poyiadjis et al. [2005] was shown in numerical examples to be stable. The method of Poyiadjis et al. [2005] is a non-standard particle implementation that avoids the particle path degeneracy problem at the expense of a computational complexity per time step which is quadratic in the number of particles, i.e. $\mathcal{O}(N^2)$; see Section 2 for more details. Supported by numerical examples, it was conjectured in Poyiadjis et al. [2011] that even under strong mixing assumptions, the variance of the estimate of the filter derivative computed with the $\mathcal{O}(N)$ methods increases at least linearly in time while that of the $\mathcal{O}(N^2)$ is uniformly bounded w.r.t. the time index. This conjecture is confirmed in this paper. Specifically, we analyze the $\mathcal{O}(N^2)$ implementation of Poyiadjis et al. [2005] in Section 3 and obtain results on the errors of the approximation, in particular, \mathbb{L}_p bounds and a Central Limit Theorem (CLT) are presented. We show that these \mathbb{L}_p bounds and asymptotic variances appearing in the CLT are uniformly bounded w.r.t. the time index when the state-space model satisfies certain mixing assumptions. In contrast, the asymptotic variance of the $\mathcal{O}(N)$ implementations, which is also captured through the CLT, is shown to increase linearly. To the best of our knowledge, these are the first results of this kind.

An important application of our results, which is discussed in detail in Section 4, is to the problem of estimating the parameters of the model (1.1)–(1.2) from observed data. The estimates of the model parameters are found by maximizing the likelihood function $p_{\theta}(y_{0:n})$ with respect to θ using a gradient ascent algorithm which relies on the particle approximation of the filter derivative. The results we present in Section 3 have bearing on the performance of the parameter estimation algorithm, which we illustrate with numerical examples in Section 4. The Appendix contains the proofs of the main results as well as that of some supporting auxiliary results. As a final remark, although the algorithms and theoretical results are presented for a state-space model, they may be reinterpreted for Feynman-Kac models as well.

1.1 Notation and definitions

We give some basic definitions from probability and operator semigroup theory. For a measurable space (E, \mathcal{E}) let $\mathcal{M}(E)$ denote the set of all finite signed measures and $\mathcal{P}(E)$ the set of all probability measures on E . The n -fold product space $E \times \cdots \times E$ is denoted by E^n . Let $\mathcal{B}(E)$ denote the Banach space of all bounded real-valued and measurable functions $\varphi : E \rightarrow \mathbb{R}$ equipped with the uniform norm $\|\varphi\| = \sup_{x \in E} |\varphi(x)|$. For $\nu \in \mathcal{M}(E)$ and $\varphi \in \mathcal{B}(E)$, let $\nu(\varphi) = \int \nu(dx) \varphi(x)$ be the Lebesgue integral of φ w.r.t. ν . If ν is a density w.r.t. some dominating measure dx on E then, $\nu(\varphi) = \int dx \nu(x) \varphi(x)$. We recall that a bounded integral kernel $M(x, dx')$ from a measurable space (E, \mathcal{E}) into an auxiliary measurable space (E', \mathcal{E}') is an operator $\varphi \mapsto M(\varphi)$ from $\mathcal{B}(E')$ into $\mathcal{B}(E)$ such that the functions

$$x \mapsto M(\varphi)(x) := \int_{E'} M(x, dx') \varphi(x')$$

are \mathcal{E} -measurable and bounded for any $\varphi \in \mathcal{B}(E')$. The kernel M also generates a dual operator $\nu \mapsto \nu M$ from $\mathcal{M}(E)$ into $\mathcal{M}(E')$ defined by

$$(\nu M)(\varphi) := \nu(M(\varphi)).$$

Given a pair of bounded integral operators (M_1, M_2) , we let $(M_1 M_2)$ the composition operator defined by $(M_1 M_2)(\varphi) = M_1(M_2(\varphi))$.

A Markov kernel is a positive and bounded integral operator M such that $M(1)(x) = 1$ for any $x \in E$. For $\varphi \in \mathcal{B}(E)$, let

$$\text{osc}(\varphi) = \sup_{x, x' \in E} |\varphi(x) - \varphi(x')|$$

and let

$$\text{Osc}_1(E) = \{\varphi \in \mathcal{B}(E) : \text{osc}(\varphi) \leq 1\}.$$

Let $\beta(M) \in [0, 1]$ denote the Dobrushin coefficient of the Markov kernel M which is defined by the formula [Del Moral, 2004, Prop. 4.2.1]:

$$\beta(M) := \sup \{ \text{osc}(M(\varphi)) ; \varphi \in \text{Osc}_1(E') \}.$$

If there exists a positive constant ρ such that the Markov kernel M satisfies

$$M(x, dz) \geq \rho M(x', dz) \text{ for all } x, x' \in E \text{ then } \beta(M) \leq 1 - \rho.$$

For two Markov kernels M_1, M_2 , $\beta(M_1 M_2) \leq \beta(M_1) \beta(M_2)$.

Given a positive function G on E , let $\Psi_G : \nu \in \mathcal{P}(E) \mapsto \Psi_G(\nu) \in \mathcal{P}(E)$ be the probability distribution defined by

$$\Psi_G(\nu)(dx) := \frac{\nu(dx) G(x)}{\nu(G)}$$

provided $\infty > \nu(G) > 0$. The definitions above also apply if ν is a density and M is a transition density. In this case all instances of $\nu(dx)$ should be replaced with $dx\nu(x)$ and $M(x, dx')$ by $dx' M(x, x')$ where dx and dx' is generic notation for the dominating measures.

It is convenient to introduce the following transition kernels:

$$\begin{aligned} Q_{\theta, n}(x_{n-1}, dx_n) &= g_\theta(y_{n-1} | x_{n-1}) dx_n f_\theta(x_n | x_{n-1}) = dx_n q_\theta(x_n | x_{n-1}), \quad n > 0, \\ Q_{\theta, k, n}(x_k, dx_n) &= (Q_{\theta, k+1} Q_{\theta, k+2} \cdots Q_{\theta, n})(x_k, dx_n), \quad 0 \leq k \leq n, \end{aligned}$$

with the convention that $Q_{\theta, n, n} = Id$, the identity operator. Note that $Q_{\theta, k, n}(1)(x_k)$ is the density of the law of $Y_{k:n-1}$ given $X_k = x_k$. For $0 \leq p \leq n$, define the potential function $G_{\theta, p, n}$ on \mathcal{X} to be

$$G_{\theta, p, n}(x_p) = Q_{\theta, p, n}(1)(x_p) / \eta_{\theta, p} Q_{\theta, p, n}(1). \quad (1.7)$$

Let the mapping $\Phi_{\theta,k,n} : \mathcal{P}(\mathcal{X}) \rightarrow \mathcal{P}(\mathcal{X})$, $0 \leq k \leq n$, be defined as follows

$$\Phi_{\theta,k,n}(\nu)(dx_n) = \frac{\nu Q_{\theta,k,n}(dx_n)}{\nu Q_{\theta,k,n}(1)}.$$

It follows that $\eta_{\theta,n} = \Phi_{\theta,k,n}(\eta_{\theta,k})$. For conciseness, we also write $\Phi_{\theta,n-1,n}$ as $\Phi_{\theta,n}$.

A key quantity that facilitates the recursive computation of the derivative of $\eta_{\theta,n}$ is the following collection of backward Markov transition kernels:

$$M_{\theta,n}(x_n, dx_{n-1}) = \frac{\eta_{\theta,n-1}(dx_{n-1}) q_{\theta}(x_n | x_{n-1})}{\eta_{\theta,n-1}(q_{\theta}(x_n | \cdot))}, \quad n > 0. \quad (1.8)$$

Their particle approximations are

$$M_{\theta,n}^N(x_n, dx_{n-1}) = \frac{\eta_{\theta,n-1}^N(dx_{n-1}) q_{\theta}(x_n | x_{n-1})}{\eta_{\theta,n-1}^N(q_{\theta}(x_n | \cdot))}. \quad (1.9)$$

These backward Markov kernels are convenient for computing certain conditional expectations and probability measures. In particular, for $\varphi \in \mathcal{B}(\mathcal{X}^2)$, we have

$$\mathbb{E}_{\theta} [\varphi(X_{n-1}, X_n) | y_{0:n-1}, x_n] = \int M_{\theta,n}(x_n, dx_{n-1}) \varphi(x_{n-1}, x_n),$$

and the law of $X_{0:n-1}$ given $X_n = x_n$ and $Y_{0:n-1} = y_{0:n-1}$ is $M_{\theta,n}(x_n, dx_{n-1}) \cdots M_{\theta,1}(x_1, dx_0)$.

Finally, the following two definitions are needed for the CLT of the particle approximation of the derivative of $\eta_{\theta,n}$. The bounded integral operator $D_{\theta,k,n}$ from \mathcal{X} into \mathcal{X}^{n+1} is defined for any $F_n \in \mathcal{B}(\mathcal{X}^{n+1})$ by

$$D_{\theta,k,n}(F_n)(x_k) := \int \left(\prod_{j=k}^1 M_{\theta,j}(x_j, dx_{j-1}) \right) \left(\prod_{j=k}^{n-1} Q_{\theta,j+1}(x_j, dx_{j+1}) \right) F_n(x_{0:n}), \quad 0 \leq k \leq n, \quad (1.10)$$

with the convention that $\prod \emptyset = 1$. The particle approximation, $D_{\theta,k,n}^N$, is defined to be

$$D_{\theta,k,n}^N(F_n)(x_k) := \int \left(\prod_{j=k}^1 M_{\theta,j}^N(x_j, dx_{j-1}) \right) \left(\prod_{j=k}^{n-1} Q_{\theta,j+1}(x_j, dx_{j+1}) \right) F_n(x_{0:n}). \quad (1.11)$$

To be concise we write

$$\eta_{\theta,k}(dx_k) D_{\theta,k,n}(x_k, dx_{0:k-1}, dx_{k+1:n}) \quad \text{as} \quad \eta_{\theta,k} D_{\theta,k,n}(dx_{0:n}).$$

(And similarly for the particle versions.) Although convention dictates that $\eta_{\theta,k} D_{\theta,k,n}$ should be understood as the measure $(\eta_{\theta,k} D_{\theta,k,n})(dx_{0:k-1}, dx_{k+1:n})$, when we mean otherwise it should be clear from the infinitesimal neighborhood.

2 Computing the filter derivative

For any $F_n \in \mathcal{B}(\mathcal{X}^{n+1})$, we have

$$\begin{aligned}
& \nabla \mathbb{Q}_{\theta,n}(F_n) \\
&= \frac{1}{p_\theta(y_{0:n-1})} \int dx_{0:n} \nabla \left(\pi_\theta(x_0) \prod_{k=1}^n f_\theta(x_k | x_{k-1}) \prod_{k=0}^{n-1} g_\theta(y_k | x_k) \right) F_n(x_{0:n}) \\
&\quad - \frac{1}{p_\theta(y_{0:n-1})} \mathbb{E}_\theta \{ F_n(X_{0:n}) | y_{0:n-1} \} \int dx_{0:n} \nabla \left(\pi_\theta(x_0) \prod_{k=1}^n f_\theta(x_k | x_{k-1}) \prod_{k=0}^{n-1} g_\theta(y_k | x_k) \right) \\
&= \mathbb{E}_\theta \{ F_n(X_{0:n}) T_{\theta,n}(X_{0:n}) | y_{0:n-1} \} - \mathbb{E}_\theta \{ F_n(X_{0:n}) | y_{0:n-1} \} \mathbb{E}_\theta \{ T_{\theta,n}(X_{0:n}) | y_{0:n-1} \} \tag{2.1}
\end{aligned}$$

where

$$T_{\theta,n}(x_{0:n}) = \sum_{k=0}^n t_{\theta,k}(x_{k-1}, x_k) \tag{2.2}$$

$$t_{\theta,k}(x_{k-1}, x_k) = \nabla \log (g_\theta(y_{k-1} | x_{k-1}) f_\theta(x_k | x_{k-1})), \quad k > 0, \tag{2.3}$$

$$t_{\theta,0}(x_{-1}, x_0) = t_{\theta,0}(x_0) = \nabla \log \pi_\theta(x_0). \tag{2.4}$$

The first equality in (2.1) follows from the definition of $\mathbb{Q}_{\theta,n}$ and interchanging the order of differentiation and integration. The interchange is permissible under certain regularity conditions [Pflug, 1996]; e.g. a sufficient condition would be the main assumption in Section 3 under which the uniform stability results are proved. The second equality follows from a change of measure, which then permits an importance sampling based estimator for the derivative of $\mathbb{Q}_{\theta,n}$; this is the well known score method, e.g. see Pflug [1996, Section 4.2.1]. For any $\varphi_n \in \mathcal{B}(\mathcal{X})$, it follows by setting $F_n(x_{0:n}) = \varphi_n(x_n)$ in (2.1) that

$$\begin{aligned}
& \nabla \int \eta_{\theta,n}(dx_n) \varphi_n(x_n) \\
&= \mathbb{E}_\theta \{ \varphi_n(X_n) T_{\theta,n}(X_{0:n}) | y_{0:n-1} \} - \mathbb{E}_\theta \{ \varphi_n(X_n) | y_{0:n-1} \} \mathbb{E}_\theta \{ T_{\theta,n}(X_{0:n}) | y_{0:n-1} \} \\
&= \int \zeta_{\theta,n}(dx_n) \varphi_n(x_n)
\end{aligned}$$

where

$$\zeta_{\theta,n}(dx_n) = \eta_{\theta,n}(dx_n) (\mathbb{E}_\theta [T_{\theta,n}(X_{0:n}) | y_{0:n-1}, x_n] - \mathbb{E}_\theta [T_{\theta,n}(X_{0:n}) | y_{0:n-1}]). \tag{2.5}$$

We call $\zeta_{\theta,n}$ the derivative of $\eta_{\theta,n}$.

Given the particle approximation (1.5) of $\mathbb{Q}_{\theta,n}$, it is straightforward to construct a particle approximation of $\zeta_{\theta,n}$:

$$\zeta_{\theta,n}^{p,N}(dx_n) = \sum_{i=1}^N \frac{1}{N} \left(T_{\theta,n}(X_{0:n}^{(i)}) - \frac{1}{N} \sum_{j=1}^N T_{\theta,n}(X_{0:n}^{(j)}) \right) \delta_{X_n^{(i)}}(dx_n). \tag{2.6}$$

This approximation is also referred to as the path space method. Such approximations were implicitly proposed in Cérou et al. [2001] and Doucet and Tadić [2003] and there are several reasons why this estimate appears attractive. Firstly, even with the resampling steps in the construction of $\mathbb{Q}_{\theta,n}^{p,N}$, $\zeta_{\theta,n}^{p,N}$ can be computed recursively. Secondly, there is no need to store the entire ancestry of each particle, i.e. $\{X_{0:n}^{(i)}\}_{1 \leq i \leq N}$, and thus the memory requirement to construct $\zeta_{\theta,n}^{p,N}$ is constant over

time. Thirdly, the computational cost per time is $\mathcal{O}(N)$. However, as $\mathbb{Q}_{\theta,n}^{\text{p},N}$ suffers from the particle path degeneracy problem, we expect the approximation $\zeta_{\theta,n}^{\text{p},N}$ to worsen over time. This was indeed observed in numerical examples in Poyiadjis et al. [2011] and it was conjectured that the asymptotic variance (i.e. as $N \rightarrow \infty$) of $\zeta_{\theta,n}^{\text{p},N}$ for bounded integrands would increase linearly with n even under strong mixing assumptions. This is now proven in this article.

An alternative particle method to approximate $\{\zeta_{\theta,n}\}_{n \geq 0}$ has been proposed in Poyiadjis et al. [2005, 2011]. We now reinterpret this method using the representation in (2.5) and a different particle approximation of $\mathbb{Q}_{\theta,n}$ that avoids the path degeneracy problem.

The measure $\mathbb{Q}_{\theta,n}$ admits the following backward representation

$$\mathbb{Q}_{\theta,n}(dx_{0:n}) = \eta_{\theta,n}(dx_n) \prod_{k=n}^1 M_{\theta,k}(x_k, dx_{k-1})$$

and the corresponding particle approximation of $\mathbb{Q}_{\theta,n}$ is given by

$$\mathbb{Q}_{\theta,n}^N(dx_{0:n}) = \eta_{\theta,n}^N(dx_n) \prod_{k=n}^1 M_{\theta,k}^N(x_k, dx_{k-1})$$

where $M_{\theta,k}^N$ was defined in (1.9). This now gives rise to the following particle approximation of $\zeta_{\theta,n}$ [Poyiadjis et al., 2005, 2011]:

$$\zeta_{\theta,n}^N(\varphi_n) = \int \mathbb{Q}_{\theta,n}^N(dx_{0:n}) T_{\theta,n}(x_{0:n}) (\varphi_n(x_n) - \eta_{\theta,n}^N(\varphi_n))$$

and indeed $\eta_{\theta,n}^N(\varphi_n) = \int \mathbb{Q}_{\theta,n}^N(dx_{0:n}) \varphi_n(x_n)$. It is apparent that $\mathbb{Q}_{\theta,n}^N$ constructed using this backward method avoids the degeneracy in paths. It is even possible to compute $\zeta_{\theta,n}^N$ recursively as detailed in Algorithm 1; since a recursion for $\eta_{\theta,n}$ is already available, it is apparent from (2.5) that what remains is to specify a recursion for $\mathbb{E}_{\theta} [T_{\theta,n}(X_{0:n}) | y_{0:n-1}, x_n]$. Let $\bar{T}_{\theta,n}(x_n)$ denote this term, then for $n \geq 1$,

$$\begin{aligned} \bar{T}_{\theta,n}(x_n) &= \mathbb{E}_{\theta} [T_{\theta,n}(X_{0:n}) | y_{0:n-1}, x_n] \\ &= \mathbb{E}_{\theta} [T_{\theta,n-1}(X_{0:n-1}) | y_{0:n-1}, x_n] + \mathbb{E}_{\theta} [t_{\theta,n}(X_{n-1}, X_n) | y_{0:n-1}, x_n] \\ &= \int M_{\theta,n}(x_n, dx_{n-1}) (\mathbb{E}_{\theta} [T_{\theta,n-1}(X_{0:n-1}) | y_{0:n-2}, x_{n-1}] + t_{\theta,n}(x_{n-1}, x_n)) \\ &= \int M_{\theta,n}(x_n, dx_{n-1}) (\bar{T}_{\theta,n-1}(x_{n-1}) + t_{\theta,n}(x_{n-1}, x_n)) \end{aligned}$$

where $\bar{T}_{\theta,0}(x_0) = t_{\theta,0}(x_0)$. Algorithm 1 computes $\zeta_{\theta,n}^N$ recursively in time by computing $(\bar{T}_{\theta,n}, \eta_{\theta,n})$ and is initialized with $\bar{T}_{\theta,0}^{(i)} = t_{\theta,0}(X_0^{(i)})$ (see (2.2)) where $\{X_0^{(i)}\}_{1 \leq i \leq N}$ are samples from $\pi_{\theta}(x_0)$.

Algorithm 1: A Particle Method to Compute the Filter Derivative

- Assume at time $n-1$ that approximate samples $\{X_{n-1}^{(i)}\}_{1 \leq i \leq N}$ from $\eta_{\theta,n-1}$ and approximations $\{\bar{T}_{\theta,n-1}^{(i)}\}_{1 \leq i \leq N}$ of $\{\bar{T}_{\theta,n-1}(X_{n-1}^{(i)})\}_{1 \leq i \leq N}$ are available.
- At time n , sample $\{X_n^{(i)}\}_{1 \leq i \leq N}$ independently from the mixture

$$\frac{\sum_{j=1}^N f_{\theta}(x_n | X_{n-1}^{(j)}) g_{\theta}(y_{n-1} | X_{n-1}^{(j)})}{\sum_{j=1}^N g_{\theta}(y_{n-1} | X_{n-1}^{(j)})} \quad (2.7)$$

and then compute $\{\bar{T}_{\theta,n}^{(i)}\}_{1 \leq i \leq N}$ and $\zeta_{\theta,n}^N$ as follows:

$$\bar{T}_{\theta,n}^{(i)} = \frac{\sum_{j=1}^N \left(\bar{T}_{\theta,n-1}^{(j)} + t_{\theta,n} \left(X_{n-1}^{(j)}, X_n^{(i)} \right) \right) f_{\theta} \left(X_n^{(i)} | X_{n-1}^{(j)} \right) g_{\theta} \left(y_{n-1} | X_{n-1}^{(j)} \right)}{\sum_{j=1}^N f_{\theta} \left(X_n^{(i)} | X_{n-1}^{(j)} \right) g_{\theta} \left(y_{n-1} | X_{n-1}^{(j)} \right)}, \quad (2.8)$$

$$\zeta_{\theta,n}^N(dx_n) = \frac{1}{N} \sum_{i=1}^N \left(\bar{T}_{\theta,n}^{(i)} - \frac{1}{N} \sum_{j=1}^N \bar{T}_{\theta,n}^{(j)} \right) \delta_{X_n^{(i)}}(dx_n). \quad (2.9)$$

Algorithm 1 uses the bootstrap particle filter of Gordon et al. [1993]. Note that any SMC implementation of $\{\eta_{\theta,n}\}_{n \geq 0}$ may be used, e.g. the auxiliary SMC method of Pitt and Shephard [1999] or sequential importance resampling with a tailored proposal distribution [Doucet et al., 2001]. It was conjectured in Poyiadjis et al. [2011] that the asymptotic variance of $\zeta_{\theta,n}^N(\varphi)$ for bounded integrands φ is uniformly bounded w.r.t. n under mixing assumptions. This is established in this article.

3 Stability of the particle estimates

The convergence analysis of $\zeta_{\theta,n}^N$ (and $\zeta_{\theta,n}^{p,N}$ for performance comparison) will largely focus on the convergence analysis of the N -particle measures $\mathbb{Q}_{\theta,n}^N$ (and correspondingly $\mathbb{Q}_{\theta,n}^{p,N}$) towards their limiting values $\mathbb{Q}_{\theta,n}$, as $N \rightarrow \infty$, which is in turn intimately related to the convergence of the flow of particle measures $\{\eta_{\theta,n}^N\}_{n \geq 0}$ towards their limiting measures $\{\eta_{\theta,n}\}_{n \geq 0}$. The \mathbb{L}_r error bounds and the central limit theorem presented here have been derived using the techniques developed in Del Moral [2004] for the convergence analysis of the particle occupation measures $\eta_{\theta,n}^N$. One of the central objects in this analysis is the local sampling errors defined as

$$V_{\theta,n}^N = \sqrt{N} (\eta_{\theta,n}^N - \Phi_{\theta,n}(\eta_{\theta,n-1}^N)) \quad (3.1)$$

The fluctuation and the deviations of these centered random measures can be estimated using non-asymptotic Kintchine's type \mathbb{L}_r -inequalities, as well as Hoeffding's or Bernstein's type exponential deviations [Del Moral, 2004, Del Moral and Rio, 2009]. In Del Moral and Miclo [2000] it is proved that these random perturbations behave asymptotically as Gaussian random perturbations; see Lemma 7.10 in the Appendix for more details. In the proof of Theorem 7.11 (a supporting theorem) in the Appendix we provide some key decompositions expressing the deviation of the particle measures $\mathbb{Q}_{\theta,n}^N$ around its limiting value $\mathbb{Q}_{\theta,n}$ in terms of the local sampling errors $(V_{\theta,0}^N, \dots, V_{\theta,n}^N)$. These decompositions are key to deriving the \mathbb{L}_r -mean error bounds and central limit theorems for the filter derivative.

The following regularity conditions are assumed.

(A) The dominating measures dx on \mathcal{X} and dy on \mathcal{Y} are finite, and there exist constants $0 < \rho, \delta, c < \infty$ such that for all $(x, x', y, \theta) \in \mathcal{X}^2 \times \mathcal{Y} \times \Theta$, the derivatives of $\pi_{\theta}(x)$, $f_{\theta}(x'|x)$ and $g_{\theta}(y|x)$ with respect to θ exists and

$$\rho^{-1} \leq f_{\theta}(x'|x) \leq \rho, \quad \delta^{-1} \leq g_{\theta}(y|x) \leq \delta, \quad (3.2)$$

$$|\nabla \log \pi_{\theta}(x)| \vee |\nabla \log f_{\theta}(x'|x)| \vee |\nabla \log g_{\theta}(y|x)| \leq c. \quad (3.3)$$

Admittedly, these conditions are restrictive and fail to hold for many models in practice. (Exceptions would include applications with a compact state-space.) However, they are typically made to establish the time uniform stability of particle approximations of the filter [Del Moral, 2004, Cappé et al., 2005] as they lead to simpler and more transparent proofs. Also, we observe that the behaviors predicted by the Theorems below seem to hold in practice even in cases where the state-space models

do not satisfy these assumptions; see Section 4. Thus the results in this paper can be seen to provide a qualitative guide to the behavior of the particle approximation even in the more general setting.

For each parameter vector $\theta \in \Theta$, realization of observations $y = \{y_n\}_{n \geq 0}$ and particle number N , let $(\Omega, \mathcal{F}, \mathbb{P}_\theta^y)$ be the underlying probability space of the random process $\{(X_n^{(1)}, \dots, X_n^{(N)})\}_{n \geq 0}$ comprised of the particle system only. Let \mathbb{E}_θ^y the corresponding expectation operator computed with respect to \mathbb{P}_θ^y . The first of the two main results in this section is a time uniform non-asymptotic error bound.

Theorem 3.1 *Assume (A). For any $r \geq 1$, there exists a constant C_r such that for all $\theta \in \Theta$, $y = \{y_n\}_{n \geq 0}$, $n \geq 0$, $N \geq 1$, and $\varphi_n \in \text{Osc}_1(\mathcal{X})$,*

$$\sqrt{N} \mathbb{E}_\theta^y \left\{ \left| \zeta_{\theta,n}^N(\varphi_n) - \zeta_{\theta,n}(\varphi_n) \right|^r \right\}^{\frac{1}{r}} \leq C_r$$

Let $\{V_{\theta,n}\}_{n \geq 0}$ be a sequence of independent centered Gaussian random fields defined as follows. For any sequence $\{\varphi_n\}_{n \geq 0}$ in $\mathcal{B}(\mathcal{X})$ and any $p \geq 0$, $\{V_{\theta,n}(\varphi_n)\}_{n=0}^p$ is a collection of independent zero-mean Gaussian random variables with variances given by

$$\eta_{\theta,n}(\varphi_n^2) - \eta_{\theta,n}(\varphi_n)^2. \quad (3.4)$$

Theorem 3.2 *Assume (A). There exists a constant $C < \infty$ such that for any $\theta \in \Theta$, $y = \{y_n\}_{n \geq 0}$, $n \geq 0$ and $\varphi_n \in \text{Osc}_1(\mathcal{X})$, $\sqrt{N} \left(\zeta_{\theta,n}^N - \zeta_{\theta,n} \right) (\varphi_n)$ converges in law, as $N \rightarrow \infty$, to the centered Gaussian random variable*

$$\sum_{p=0}^n V_{\theta,p} \left(G_{\theta,p,n} \frac{D_{\theta,p,n}(F_{\theta,n} - \mathbb{Q}_{\theta,n}(F_{\theta,n}))}{D_{\theta,p,n}(1)} \right) \quad (3.5)$$

whose variance is uniformly bounded above by C where

$$F_{\theta,n} = (\varphi_n - \mathbb{Q}_{\theta,n}(\varphi_n)) (T_{\theta,n} - \mathbb{Q}_{\theta,n}(T_{\theta,n})).$$

The proofs of both these results are in the Appendix.

As a comparison, we quantify the variance of the particle estimate of the filter derivative computed using the path-based method (see (2.6).) Consider the following simplified example that serves to illustrate the point. Let $g_\theta(y|x) = g(y|x)$ (that is θ -independent), $f_\theta(x_n|x_{n-1}) = \pi_\theta(x_n)$, where π_θ is the initial distribution. (Note that f_θ in this case satisfies a rephrased version of (3.2) under which the conclusion of Theorem 3.2 also holds.) Also, consider the sequence of repeated observations $y_0 = y_1 = \dots$ where y_0 is arbitrary. Applying Lemma 7.12 (in the Appendix) that characterizes the limiting distribution of $\sqrt{N}(\mathbb{Q}_{\theta,n}^{p,N} - \mathbb{Q}_{\theta,n})$ to this special case results in $\sqrt{N}(\zeta_{\theta,n}^{p,N} - \zeta_{\theta,n})(\varphi)$ (see (2.6)) having an asymptotic distribution which is Gaussian with mean zero and variance

$$n \times \pi_\theta(\bar{\varphi}^2) \pi'_\theta [(\nabla \log \pi_\theta)^2] + \pi_\theta [\bar{\varphi}^2 (\nabla \log \pi_\theta)^2] - \nabla \pi_\theta(\varphi)^2$$

where $\bar{\varphi} = \varphi - \pi_\theta(\varphi)$, $\pi'_\theta(x) = \pi_\theta(x)g(y_0|x)/\pi_\theta(g(y_0|\cdot))$. This variance increases linearly with time in contrast to the time bounded variance of Theorem 3.2.

4 Application to recursive parameter estimation

Being able to compute $\{\zeta_{\theta,n}\}_{n \geq 0}$ is particularly useful when performing online static parameter estimation for state-space models using Recursive Maximum Likelihood (RML) techniques [Le Gland and Mevel, 1997, Poyiadjis et al., 2005, 2011]; see also Kantas et al. [2009] for a general review of available particle methods based solutions, including Bayesian ones, for this problem. The computed filter derivative may also be useful in other areas; e.g. see Coquelin et al. [2008] for an application in control.

4.1 Recursive Maximum Likelihood

Let θ^* be the true static parameter generating the observed data $\{y_n\}_{n \geq 0}$. Given a finite record of observations $y_{0:T}$, the log-likelihood may be maximized with the following steepest ascent algorithm:

$$\theta_k = \theta_{k-1} + \gamma_k \nabla \log p_\theta(y_{0:T})|_{\theta=\theta_{k-1}}, \quad k \geq 1, \quad (4.1)$$

where θ_0 is some arbitrary initial guess of θ^* , $\nabla \log p_\theta(y_{0:T})|_{\theta=\theta_{k-1}}$ denotes the gradient of the log-likelihood evaluated at the current parameter estimate and $\{\gamma_k\}_{k \geq 1}$ is a decreasing positive real-valued step-size sequence, which should satisfy the following constraints:

$$\sum_{k=1}^{\infty} \gamma_k = \infty, \quad \sum_{k=1}^{\infty} \gamma_k^2 < \infty.$$

Although $\nabla \log p_\theta(y_{0:T})$ can be computed using (4.3), the computation cost can be prohibitive for a long data record since each iteration of (4.1) would require a complete browse through the $T + 1$ data points. A more attractive alternative would be a recursive procedure in which the data is run through once only sequentially. For example, consider the following update scheme:

$$\theta_n = \theta_{n-1} + \gamma_n \nabla \log p_\theta(y_n|y_{0:n-1})|_{\theta=\theta_{n-1}} \quad (4.2)$$

where $\nabla \log p_\theta(y_n|y_{0:n-1})|_{\theta=\theta_{n-1}}$ denotes the gradient of $\log p_\theta(y_n|y_{0:n-1})$ evaluated at the current parameter estimate; that is upon receiving y_n , θ_{n-1} is updated in the direction of ascent of the conditional density of this new observation. Since we have

$$\nabla \log p_\theta(y_n|y_{0:n-1})|_{\theta=\theta_{n-1}} = \frac{\int dx_n \eta_{\theta_{n-1},n}(x_n) \nabla g_\theta(y_n|x_n)|_{\theta=\theta_{n-1}} + \int dx_n (y_n|x_n) \zeta_{\theta_{n-1},n}(x_n) g_{\theta_{n-1}}}{\int dx_n \eta_{\theta_{n-1},n}(x_n) g_{\theta_{n-1}}(y_n|x_n)}, \quad (4.3)$$

this clearly requires the filter derivative $\zeta_{\theta,n}$. The algorithm in the present form is not suitable for online implementation as it requires re-computing the filter and its derivative at the value $\theta = \theta_{n-1}$ from time zero. The RML procedure uses an approximation of (4.3) which is obtained by updating the filter and its derivative using the parameter value θ_{n-1} at time n ; we refer the reader to Le Gland and Mevel [1997] for details. The asymptotic properties of the RML algorithm, i.e. the behavior of θ_n in the limit as n goes to infinity, has been studied in the case of an i.i.d. hidden process by Titterton [1984] and Le Gland and Mevel [1997] for a finite state-space hidden Markov model. It is shown in Le Gland and Mevel [1997] that under regularity conditions this algorithm converges towards a local maximum of the average log-likelihood and that this average log-likelihood is maximized at θ^* . A particle version of the RML algorithm of Le Gland and Mevel [1997] that uses Algorithm 1's estimate of $\eta_{\theta,n}$ is presented as Algorithm 2.

Algorithm 2: Particle Recursive Maximum Likelihood

- At time $n - 1$ we are given $y_{0:n-1}$, the previous estimate θ_{n-1} of θ^* and $\{(X_{n-1}^{(i)}, \bar{T}_{n-1}^{(i)})\}_{i=1}^N$.
- At time n , upon receiving y_n , sample $\{X_n^{(i)}\}_{1 \leq i \leq N}$ independently from (2.7) using parameter $\theta = \theta_{n-1}$ to obtain

$$\eta_n^N(dx_n) = \frac{1}{N} \sum_{i=1}^N \delta_{X_n^{(i)}}(dx_n)$$

and then compute

$$\bar{T}_n^{(i)} = \frac{\sum_{j=1}^N \left(\bar{T}_{n-1}^{(j)} + t_{\theta_{n-1}, n} \left(X_{n-1}^{(j)}, X_n^{(i)} \right) \right) f_{\theta_{n-1}} \left(X_n^{(i)} | X_{n-1}^{(j)} \right) g_{\theta_{n-1}} \left(y_{n-1} | X_{n-1}^{(j)} \right)}{\sum_{j=1}^N f_{\theta_{n-1}} \left(X_n^{(i)} | X_{n-1}^{(j)} \right) g_{\theta_{n-1}} \left(y_{n-1} | X_{n-1}^{(j)} \right)}, \quad (4.4)$$

$$\zeta_n^N(dx_n) = \frac{1}{N} \sum_{i=1}^N \left(\bar{T}_n^{(i)} - \frac{1}{N} \sum_{j=1}^N \bar{T}_n^{(j)} \right) \delta_{X_n^{(i)}}(dx_n), \quad (4.5)$$

and

$$\widehat{\nabla} \log p(y_n | y_{0:n-1}) = \frac{\int \eta_n^N(dx_n) \nabla g_{\theta}(y_n | x_n)|_{\theta_{n-1}} + \int \zeta_n^N(dx_n) g_{\theta_{n-1}}(y_n | x_n)}{\int \eta_n^N(dx_n) g_{\theta_{n-1}}(y_n | x_n)}.$$

Finally update the parameter:

$$\theta_n = \theta_{n-1} + \gamma_n \widehat{\nabla} \log p(y_n | y_{0:n-1}). \quad (4.6)$$

Under Assumption A, the particle approximation of the filter is stable [Del Moral, 2004]; see also Lemma 7.4 in the Appendix. This combined with the proven stability of the particle approximation of the filter derivative implies that the particle estimate of the derivative of $\log p(y_n | y_{0:n-1})$ is also stable.

4.2 Simulations

The RML algorithm is applied to the following stochastic volatility model [Pitt and Shephard, 1999]:

$$\begin{aligned} X_0 &\sim \mathcal{N}\left(0, \frac{\sigma^2}{1 - \phi^2}\right), \quad X_{n+1} = \phi X_n + \sigma V_{n+1}, \\ Y_n &= \beta \exp(X_n/2) W_n, \end{aligned}$$

where $\mathcal{N}(m, s)$ denotes a Gaussian random variable with mean m and variance s , $V_n \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, 1)$ and $W_n \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, 1)$ are two mutually independent sequences, both independent of the initial state X_0 . The model parameters, $\theta = (\phi, \sigma, \beta)$, are to be estimated.

Our first example demonstrates the theoretical results in Section 3. The estimate of $\partial/\partial\sigma \log p(y_{n:n+L-1} | y_{0:n-1})$ at $\theta^* = (0.8, \sqrt{0.1}, 1)$ was computed using Algorithm 1 with 500 particles and using the path-space method (see (2.6)) with 2.5×10^5 particles for the stochastic volatility model. The block size L was 500. Shown in Figure 1 is the variance of these particle estimates for various values of n derived from many independent random replications of the simulation. The linear increase of the variance of the path-space method as predicted by theory is evident although Assumption A is not satisfied.

For the path-space method, because the variance of the estimate of the filter derivative grows linearly in time, the eventual high variance in the gradient estimate can result in the divergence of the parameter estimates. To illustrate this point, (4.6) was implemented with the path-space estimate of the filter derivative (2.6) computed with 10000 particles and constant step-size sequence, $\gamma_n = 10^{-4}$ for all n . θ_0 was initialized at the true parameter value. A sequence of two million observations was simulated with $\theta^* = (0.8, \sqrt{0.1}, 1)$. The results are shown in Figure 3.

For the same value of θ^* and sequence of observations used in the previous example, Algorithm 2 was executed with 500 particles and $\gamma_n = 0.01$, $n \leq 10^5$, $\gamma_n = (n - 5 \times 10^4)^{-0.6}$, $n > 10^5$. As it




Figure 1: Variance of the particle estimates of $\partial/\partial\sigma \log p(y_{n:n+500-1}|y_{0:n-1})$ for various values of n for the stochastic volatility model. Circles are variance of Algorithm 1's estimate with 500 particles. Stars indicate the variance of the estimate of the path-space method with 2.5×10^5 particles. Dotted line is best fitting straight line to path-space method's variance to indicate trend.




Figure 2: Sequence of recursive parameter estimates, $\theta_n = (\sigma_n, \phi_n, \beta_n)$, computed using (4.6) with $N = 500$. From top to bottom: β_n , ϕ_n and σ_n and marked on the right are the “converged values” which were taken to be the empirical average of the last 1000 values.




Figure 3: RML for stochastic volatility with path-space gradient estimate with 10,000 particles, constant step-size and initialized at the true parameter values which are indicated by the dashed lines. From top to bottom, ϕ , β and σ .

can be seen from the results in Figure 2 the estimate converges to a value in the neighborhood of the true parameter.

5 Conclusion

We have presented theoretical results establishing the uniform stability of the particle approximation of the optimal filter derivative proposed in Poyiadjis et al. [2005, 2009]. While these results have been presented in the context of state-space models, they can also be applied to Feynman-Kac models [Del Moral, 2004] which could potentially enlarge the range of applications. For example, if $dx' f_\theta(x'|x)$ is reversible w.r.t. to some probability measure μ_θ and if we replace $g_\theta(y_n|x_n)$ with a time-homogeneous potential function $g_\theta(x_n)$ then $\eta_{\theta,n}$ converges, as $n \rightarrow \infty$, to the probability measure $\mu_{\theta,h}$ defined as

$$\mu_{\theta,h}(dx) := \frac{1}{\mu_\theta(h_\theta \int dx' f_\theta(x'|\cdot) h_\theta(x'))} \mu_\theta(dx) h_\theta(x) \int dx' f_\theta(x'|x) h_\theta(x')$$

where h_θ is a positive eigenmeasure associated with the top eigenvalue of the integral operator $Q_\theta(x, dx') = g_\theta(x) dx' f_\theta(x'|x)$ (see section 12.4 of Del Moral [2004]). The measure $\mu_{\theta,h}$ is the invariant measure of the h -process defined as the Markov chain with transition kernel $M_\theta(x, dx') \propto dx' f_\theta(x'|x) h_\theta(x')$. The particle algorithm described here can be directly used to approximate the derivative of this invariant measure w.r.t to θ . It would also be of interest to weaken Assumption A and there are several ways this might be approached. For example for non-ergodic signals using ideas in Oudjane and Rubenthaler [2005], Heine and Crisan [2008] or via Foster-Lyapunov conditions as in Beskos et al. [2011], Whiteley [2011].

6 Acknowledgement

We are grateful to Sinan Yildirim for carefully reading this report.

7 Appendix

The statement of the results in this section hold for any θ and any sequence of observations $y = \{y_n\}_{n \geq 0}$. All mathematical expectations are taken with respect to the law of the particle system only for the specific θ and y under consideration. While θ is retained in the statement of the results, it is omitted in the proofs. The superscript y of the expectation operator is also omitted in the proofs.

This section commences with some essential definitions in addition to those in Section 1.1. Let

$$P_{\theta,k,n}(x_k, dx_n) = \frac{Q_{\theta,k,n}(x_k, dx_n)}{Q_{\theta,k,n}(1)(x_k)},$$

and

$$\mathcal{M}_{\theta,p}(x_p, dx_{0:p-1}) = \prod_{k=p}^1 M_{\theta,k}(x_k, dx_{k-1}), \quad p > 0,$$

and its corresponding particle approximation is

$$\mathcal{M}_{\theta,p}^N(x_p, dx_{0:p-1}) = \prod_{k=p}^1 M_{\theta,k}^N(x_k, dx_{k-1})$$

To make the subsequent expressions more terse, let

$$\tilde{\eta}_{\theta,n}^N = \Phi_{\theta,n}(\eta_{\theta,n-1}^N), \quad n \geq 0, \quad (7.1)$$

where $\tilde{\eta}_{\theta,0}^N = \Phi_{\theta,0}(\eta_{-1}^N) = \eta_{\theta,0} = \pi_\theta$ by convention. (Recall $\Phi_{\theta,n} = \Phi_{\theta,n-1,n}$.) Let

$$\mathcal{F}_n^N = \sigma \left(\left\{ X_k^{(i)}; 0 \leq k \leq n, 1 \leq i \leq N \right\} \right), \quad n \geq 0,$$

be the natural filtration associated with the N -particle approximation model and let \mathcal{F}_{-1}^N be the trivial sigma field.

The following estimates are a straightforward consequence of Assumption (A). For all θ and time indices $0 \leq k < q \leq n$,

$$b_{\theta,k,n} = \sup_{x_k, x'_k} \frac{Q_{\theta,k,n}(1)(x_k)}{Q_{\theta,k,n}(1)(x'_k)} \leq \rho^2 \delta^2, \quad \beta \left(\frac{Q_{\theta,k,q}(x_k, dx_q) Q_{\theta,q,n}(1)(x_q)}{Q_{\theta,k,q}(Q_{\theta,q,n}(1))(x_k)} \right) \leq (1 - \rho^{-4})^{(q-k)} = \bar{\rho}^{q-k}, \quad (7.2)$$

and for θ , $0 < k \leq q$,

$$M_{\theta,k}^N(x, dz) \leq \rho^4 M_{\theta,k}^N(x', dz) \implies \beta(M_{\theta,q}^N \cdots M_{\theta,k}^N) \leq (1 - \rho^{-4})^{q-k+1}. \quad (7.3)$$

Note that setting $q = n$ in (7.2) yields an estimate for $\beta(P_{\theta,k,n})$

Several auxiliary results are now presented, all of which hinge on the following Kintchine type moment bound proved in Del Moral [2004, Lem. 7.3.3].

Lemma 7.1 *Del Moral [2004, Lemma 7.3.3] Let μ be a probability measure on the measurable space (E, \mathcal{E}) . Let G and h be \mathcal{E} -measurable functions satisfying $G(x) \geq cG(x') > 0$ for all $x, x' \in E$ where c is some finite positive constant. Let $\{X^{(i)}\}_{1 \leq i \leq N}$ be a collection of independent random samples from μ . If h has finite oscillation then for any integer $r \geq 1$ there exists a finite constant a_r , independent of N , G and h , such that*

$$\sqrt{N} \mathbb{E} \left\{ \left| \frac{\sum_{i=1}^N G(X^{(i)}) h(X^{(i)})}{\sum_{i=1}^N G(X^{(i)})} - \frac{\mu(Gh)}{\mu(G)} \right|^r \right\}^{\frac{1}{r}} \leq c^{-1} \text{osc}(h) a_r.$$

Proof:

The result for $G = 1$ and $c = 1$ is proved in Del Moral [2004]. The case stated here can be established using the representation

$$\frac{\mu^N(Gh)}{\mu^N(G)} - \frac{\mu(Gh)}{\mu(G)} = \frac{\mu(G)}{\mu^N(G)} (\mu^N - \mu) \left[\frac{G}{\mu(G)} \left(h - \frac{\mu(Gh)}{\mu(G)} \right) \right]$$

where $\mu^N(dx) = N^{-1} \sum_{i=1}^N \delta_{X^{(i)}}(dx)$.

Remark 7.2 For $k \geq 0$, let h_{k-1}^N be a \mathcal{F}_{k-1}^N measurable function satisfying $h_{k-1}^N \in \text{Osc}_1(\mathcal{X})$ almost surely. Then Lemma 7.1 can be invoked to establish

$$\sqrt{N} \mathbb{E}_\theta^y \left\{ \left| \frac{\eta_{\theta,k}^N(Gh_{k-1}^N)}{\eta_{\theta,k}^N(G)} - \frac{\Phi_{\theta,k}(\eta_{\theta,k-1}^N)(Gh_{k-1}^N)}{\Phi_{\theta,k}(\eta_{\theta,k-1}^N)(G)} \right|^r \right\}^{\frac{1}{r}} \leq c^{-1} a_r$$

where G is defined as in Lemma 7.1.

Lemma 7.3 to Lemma 7.6 are a consequence of Lemma 7.1 and the estimates in (7.2).

Lemma 7.3 For any $r \geq 1$ there exist a finite constant a_r such that the following inequality holds for all $\theta, y, 0 \leq k \leq n$ and \mathcal{F}_{k-1}^N measurable function φ_n^N satisfying $\varphi_n^N \in \text{Osc}_1(\mathcal{X})$ almost surely,

$$\sqrt{N} \mathbb{E}_\theta^y \left(\left| \Phi_{\theta,k,n}(\eta_{\theta,k}^N)(\varphi_n^N) - \Phi_{\theta,k-1,n}(\eta_{\theta,k-1}^N)(\varphi_n^N) \right|^r \right)^{\frac{1}{r}} \leq a_r b_{\theta,k,n} \beta(P_{\theta,k,n}),$$

where, by convention $\Phi_{\theta,-1,n}(\eta_{\theta,-1}^N) = \eta_{\theta,n}$, and the constants $b_{\theta,k,n}$ and $\beta(P_{\theta,k,n})$ were defined in (7.2).

Proof:

$$\begin{aligned} & \Phi_{k,n}(\eta_k^N)(\varphi_n^N) - \Phi_{k-1,n}(\eta_{k-1}^N)(\varphi_n^N) \\ &= \int \left(\frac{\eta_k^N(dx_k) Q_{k,n}(1)(x_k)}{\eta_k^N Q_{k,n}(1)} - \frac{\Phi_k(\eta_{k-1}^N)(dx_k) Q_{k,n}(1)(x_k)}{\Phi_k(\eta_{k-1}^N) Q_{k,n}(1)} \right) P_{k,n}(\varphi_n^N)(x_k) \end{aligned}$$

where $\Phi_0(\eta_{-1}^N) = \eta_0$ by convention. Applying Lemma 7.1 with the estimates in (7.2) we have

$$\sqrt{N} \mathbb{E} \left(\left| \Phi_{k,n}(\eta_k^N)(\varphi_n^N) - \Phi_{k-1,n}(\eta_{k-1}^N)(\varphi_n^N) \right|^r \mid \mathcal{F}_{k-1}^N \right)^{\frac{1}{r}} \leq a_r b_{k,n} \beta(P_{k,n})$$

almost surely.

Lemma 7.3 may be used to derive the following error estimate [Del Moral, 2004, Theorem 7.4.4].

Lemma 7.4 For any $r \geq 1$, there exists a constant c_r such that the following inequality holds for all $\theta, y, n \geq 0$ and $\varphi \in \text{Osc}_1(\mathcal{X})$,

$$\sqrt{N} \mathbb{E}_\theta^y \left(\left| [\eta_{\theta,n}^N - \eta_{\theta,n}](\varphi) \right|^r \right)^{\frac{1}{r}} \leq c_r \sum_{k=0}^n b_{\theta,k,n} \beta(P_{\theta,k,n}). \quad (7.4)$$

Assume (A). For any $r \geq 1$, there exists a constant c'_r such that for all $\theta, y, n \geq 0, \varphi \in \text{Osc}_1(\mathcal{X})$, $G \in \mathcal{B}(\mathcal{X})$ such that G is positive and satisfies $G(x) \geq c_G G(x')$ for all $x, x' \in \mathcal{X}$ for some positive constant c_G ,

$$\sqrt{N} \mathbb{E}_\theta^y \left(\left| \left[\frac{\eta_{\theta,n}^N(dx_n)G(x_n)}{\eta_{\theta,n}^N(G)} - \frac{\eta_{\theta,n}(dx_n)G(x_n)}{\eta_{\theta,n}(G)} \right] (\varphi) \right|^r \right)^{\frac{1}{r}} \leq c'_r (1 + c_G^{-1}). \quad (7.5)$$

Proof:

The first part follows from applying Lemma 7.3 to the telescopic sum [Del Moral, 2004, Theorem 7.4.4]:

$$(\eta_n^N - \eta_n)(\varphi) = \sum_{k=0}^n \Phi_{k,n}(\eta_k^N)(\varphi) - \Phi_{k-1,n}(\eta_{k-1}^N)(\varphi)$$

with the convention that $\Phi_{-1,n}(\eta_{-1}^N) = \eta_n$. For the second part, use the same telescopic sum but with the k -th term being

$$\begin{aligned} & \frac{\Phi_{k,n}(\eta_k^N)(\varphi G)}{\Phi_{k,n}(\eta_k^N)(G)} - \frac{\Phi_{k-1,n}(\eta_{k-1}^N)(\varphi G)}{\Phi_{k-1,n}(\eta_{k-1}^N)(G)} \\ &= \int \left(\frac{\eta_k^N(dx_k)Q_{k,n}(G)(x_k)}{\eta_k^N Q_{k,n}(G)} - \frac{\Phi_k(\eta_{k-1}^N)(dx_k)Q_{k,n}(G)(x_k)}{\Phi_k(\eta_{k-1}^N)Q_{k,n}(G)} \right) \frac{Q_{k,n}(G\varphi)(x_k)}{Q_{k,n}(G)(x_k)}. \end{aligned}$$

Apply Lemma 7.1 using the same estimates in (7.2), i.e. the same estimates hold with G replacing 1 in the definition of $b_{k,n}$ and with G replacing $Q_{q,n}(1)$ in the argument of β .

The following result is a consequence of Lemma 7.4.

Lemma 7.5 Assume (A). For any $r \geq 1$, there exists a constant c_r such that the following inequality holds for all $\theta, y, 0 \leq k \leq n, N > 0$ and $\varphi_n \in \text{Osc}_1(\mathcal{X})$,

$$\sqrt{N} \mathbb{E}_\theta^y \left(\left| [\Phi_{\theta,k,n}(\eta_{\theta,k}^N) - \Phi_{\theta,k,n}(\eta_{\theta,k})] (\varphi_n) \right|^r \right)^{\frac{1}{r}} \leq c_r \bar{\rho}^{n-k}$$

Proof:

The result is established by expressing $\Phi_{k,n}(\eta_k^N)$ as

$$\Phi_{k,n}(\eta_k^N)(dx_n) = \int \frac{\eta_k^N(dx_k)Q_{k,n}(1)(x_k)}{\eta_k^N Q_{k,n}(1)} P_{k,n}(x_k, dx_n),$$

expressing $\Phi_{k,n}(\eta_k)$ similarly, setting G in (7.5) to $Q_{k,n}(1)$, $\varphi = P_{k,n}(\varphi_n)$ and using the estimates in (7.2).

Lemma 7.6 For each $r \geq 1$, there exists a finite constant c_r such that for all $\theta, y, 0 \leq k \leq q \leq n$, and \mathcal{F}_{k-1}^N measurable functions φ_q^N satisfying $\varphi_q \in \text{Osc}_1(\mathcal{X})$ almost surely,

$$\begin{aligned} & \sqrt{N} \mathbb{E}_\theta^y \left(\left| \int \left(\frac{\Phi_{\theta,k,q}(\eta_{\theta,k}^N)(dx_q)Q_{\theta,q,n}(1)(x_q)}{\Phi_{\theta,k,q}(\eta_{\theta,k}^N)Q_{\theta,q,n}(1)} - \frac{\Phi_{\theta,k-1,q}(\eta_{\theta,k-1}^N)(dx_q)Q_{\theta,q,n}(1)(x_q)}{\Phi_{\theta,k-1,q}(\eta_{\theta,k-1}^N)Q_{\theta,q,n}(1)} \right) \varphi_q^N(x_q) \right|^r \right)^{\frac{1}{r}} \\ & \leq c_r b_{\theta,k,n} \beta \left(\frac{Q_{\theta,k,q}(x_k, dx_q)Q_{\theta,q,n}(1)(x_q)}{Q_{\theta,k,q}(Q_{\theta,q,n}(1))(x_k)} \right) \end{aligned}$$

Proof:

This results is established by noting that

$$\begin{aligned} & \frac{\Phi_{k,q}(\eta_k^N)(dx_q)Q_{q,n}(1)(x_q)}{\Phi_{k,q}(\eta_k^N)Q_{q,n}(1)} - \frac{\Phi_{k-1,q}(\eta_{k-1}^N)(dx_q)Q_{q,n}(1)(x_q)}{\Phi_{k-1,q}(\eta_{k-1}^N)Q_{q,n}(1)} \\ &= \int \left(\frac{\eta_k^N(dx_k)Q_{k,n}(1)(x_k)}{\eta_k^N Q_{k,n}(1)} - \frac{\Phi_k(\eta_{k-1}^N)(dx_k)Q_{k,n}(1)(x_k)}{\Phi_k(\eta_{k-1}^N)Q_{k,n}(1)} \right) \frac{Q_{k,q}(x_k, dx_q)Q_{q,n}(1)(x_q)}{Q_{k,n}(1)(x_k)}. \end{aligned}$$

Now Lemma 7.1 is applied using the estimates in (7.2).

Lemma 7.7 *Assume (A). There exists a collection of a pair of finite positive constants, $a_i, c_i, i \geq 1$, such that the following bounds hold for all $r \geq 1, \theta, y, 0 \leq p \leq n, N \geq 1, x_p \in \mathcal{X}, F_p \in \mathcal{B}(\mathcal{X}^{p+1}), F_n \in \mathcal{B}(\mathcal{X}^{n+1})$,*

$$\begin{aligned} \sqrt{N} \mathbb{E}_\theta^y \left(\left| \mathcal{M}_{\theta,p}^N(F_p(\cdot, x_p))(x_p) - \mathcal{M}_{\theta,p}(F_p(\cdot, x_p))(x_p) \right|^r \right)^{\frac{1}{r}} &\leq \|F_p\| a_r p, \\ \sqrt{N} \mathbb{E}_\theta^y \left(\left| D_{\theta,p,n}^N(F_n)(x_p) - D_{\theta,p,n}(F_n)(x_p) \right|^r \right)^{\frac{1}{r}} &\leq a_r c_n \|F_n\|. \end{aligned}$$

Proof:

For each x_p , let $x_{0:p-1} \rightarrow G_{p-1,x_p}(x_{0:p-1}) = F_p(x_{0:p})q(x_p|x_{p-1})$. Adopting the convention $\tilde{\eta}_0^N = \eta_0$,

$$\begin{aligned} & \mathcal{M}_p^N(F_p(\cdot, x_p))(x_p) - \mathcal{M}_p(F_p(\cdot, x_p))(x_p) \\ &= \sum_{k=1}^p \int \left(\frac{\eta_{p-k}^N D_{p-k,p-1}^N(dx_{0:p-1})q(x_p|x_{p-1})}{\eta_{p-k}^N D_{p-k,p-1}^N(q(x_p|\cdot))} - \frac{\tilde{\eta}_{p-k}^N D_{p-k,p-1}^N(dx_{0:p-1})q(x_p|x_{p-1})}{\tilde{\eta}_{p-k}^N D_{p-k,p-1}^N(q(x_p|\cdot))} \right) F_p(x_{0:p}) \\ &= \sum_{k=1}^p \int \left(\frac{\eta_{p-k}^N(dx_{p-k})Q_{p-k,p-1}(q(x_p|\cdot))(x_{p-k})}{\eta_{p-k}^N Q_{p-k,p-1}(q(x_p|\cdot))} - \frac{\tilde{\eta}_{p-k}^N(dx_{p-k})Q_{p-k,p-1}(q(x_p|\cdot))(x_{p-k})}{\tilde{\eta}_{p-k}^N Q_{p-k,p-1}(q(x_p|\cdot))} \right) \\ &\quad \times \frac{G_{p-k,p-1,x_p}^N(x_{p-k})}{Q_{p-k,p-1}(q(x_p|\cdot))(x_{p-k})} \end{aligned}$$

where $G_{p-k,p-1,x_p}^N(x_{p-k}) = D_{p-k,p-1}^N(G_{p-1,x_p})(x_{p-k})$, which is a \mathcal{F}_{p-k-1}^N -measurable function with norm

$$\sup_{x_{p-k}} \left| \frac{G_{p-k,p-1,x_p}^N(x_{p-k})}{Q_{p-k,p-1}(q(x_p|\cdot))(x_{p-k})} \right| \leq \|F_p\|.$$

The result is established upon applying Lemma 7.1 (see Remark 7.2) to each term in the sum separately and using the estimates in (7.2). To establish the second result, let

$$F_{p,n}(x_{0:p}) = \int Q_{p+1}(x_p, dx_{p+1}) \cdots Q_n(x_{n-1}, dx_n) F_n(x_{0:n}).$$

Then,

$$D_{p,n}^N(F_n)(x_p) - D_{p,n}(F_n)(x_p) = \mathcal{M}_p^N(F_{p,n}(\cdot, x_p))(x_p) - \mathcal{M}_p(F_{p,n}(\cdot, x_p))(x_p).$$

The result follows by setting $c_n = p \sup_\theta \|Q_{\theta,p,n}(1)\|$ and it follows from Assumption (A) that c_n is finite.

Lemma 7.8 and Lemma 7.9 both build on the previous results and are needed for the proof of Theorem 3.1.

Lemma 7.8 Assume (A). For any $r \geq 1$ there exists a constant C_r such that for all $\theta, y, 0 \leq k < n$, $N \geq 1$, $\varphi_n \in \text{Osc}_1(\mathcal{X})$,

$$\begin{aligned} & \sqrt{N} \mathbb{E}_\theta^y \left\{ \left| \int \mathbb{Q}_{\theta,n}^N(dx_{0:n}) t_{\theta,k}(x_{k-1}, x_k) (\varphi_n(x_n) - \eta_{\theta,n}^N(\varphi_n)) \right. \right. \\ & \quad \left. \left. - \int \frac{\eta_{\theta,k}^N D_{\theta,k,n}^N(dx_{0:n})}{\eta_{\theta,k}^N D_{\theta,k,n}^N(1)} t_{\theta,k}(x_{k-1}, x_k) \left(\varphi_n(x_n) - \frac{\eta_{\theta,k}^N D_{\theta,k,n}^N(\varphi_n)}{\eta_{\theta,k}^N D_{\theta,k,n}^N(1)} \right) \right|^r \right\}^{\frac{1}{r}} \\ & \leq 2(n-k) C_r \bar{\rho}^{n-k} \end{aligned} \quad (7.6)$$

Proof:

The term (7.6) can be further expanded as

$$\begin{aligned} & \int \frac{\eta_k^N D_{k,n}^N(dx_{0:n})}{\eta_k^N D_{k,n}^N(1)} t_k(x_{k-1}, x_k) \left(\varphi_n(x_n) - \frac{\eta_k^N D_{k,n}^N(\varphi_n)}{\eta_k^N D_{k,n}^N(1)} \right) \\ & - \int \mathbb{Q}_n^N(dx_{0:n}) t_k(x_{k-1}, x_k) (\varphi_n(x_n) - \eta_n^N(\varphi_n)) \\ & = \sum_{p=k}^{n-1} \int \frac{\eta_p^N D_{p,n}^N(dx_{0:n})}{\eta_p^N D_{p,n}^N(1)} t_k(x_{k-1}, x_k) \left(\varphi_n(x_n) - \frac{\eta_p^N D_{p,n}^N(\varphi_n)}{\eta_p^N D_{p,n}^N(1)} \right) \\ & - \sum_{p=k}^{n-1} \int \frac{\eta_{p+1}^N D_{p+1,n}^N(dx_{0:n})}{\eta_{p+1}^N D_{p+1,n}^N(1)} t_k(x_{k-1}, x_k) \left(\varphi_n(x_n) - \frac{\eta_{p+1}^N D_{p+1,n}^N(\varphi_n)}{\eta_{p+1}^N D_{p+1,n}^N(1)} \right) \\ & = \sum_{p=k}^{n-1} \int \left(\frac{\eta_p^N D_{p,n}^N(dx_{0:n})}{\eta_p^N D_{p,n}^N(1)} - \frac{\eta_{p+1}^N D_{p+1,n}^N(dx_{0:n})}{\eta_{p+1}^N D_{p+1,n}^N(1)} \right) t_k(x_{k-1}, x_k) \left(\varphi_n(x_n) - \frac{\eta_p^N D_{p,n}^N(\varphi_n)}{\eta_p^N D_{p,n}^N(1)} \right) \\ & - \sum_{p=k}^{n-1} \left(\frac{\eta_p^N D_{p,n}^N(\varphi_n)}{\eta_p^N D_{p,n}^N(1)} - \frac{\eta_{p+1}^N D_{p+1,n}^N(\varphi_n)}{\eta_{p+1}^N D_{p+1,n}^N(1)} \right) \left(\frac{\eta_{p+1}^N D_{p+1,n}^N(t_k)}{\eta_{p+1}^N D_{p+1,n}^N(1)} - \frac{\eta_p^N D_{p,n}^N(t_k)}{\eta_p^N D_{p,n}^N(1)} \right) \\ & - \sum_{p=k}^{n-1} \left(\frac{\eta_p^N D_{p,n}^N(\varphi_n)}{\eta_p^N D_{p,n}^N(1)} - \frac{\eta_{p+1}^N D_{p+1,n}^N(\varphi_n)}{\eta_{p+1}^N D_{p+1,n}^N(1)} \right) \frac{\eta_p^N D_{p,n}^N(t_k)}{\eta_p^N D_{p,n}^N(1)} \\ & = \sum_{p=k}^{n-1} \int \left(\frac{\eta_p^N D_{p,n}^N(dx_{0:n})}{\eta_p^N D_{p,n}^N(1)} - \frac{\eta_{p+1}^N D_{p+1,n}^N(dx_{0:n})}{\eta_{p+1}^N D_{p+1,n}^N(1)} \right) \\ & \quad \times \left(t_k(x_{k-1}, x_k) - \frac{\eta_p^N D_{p,n}^N(t_k)}{\eta_p^N D_{p,n}^N(1)} \right) \left(\varphi_n(x_n) - \frac{\eta_p^N D_{p,n}^N(\varphi_n)}{\eta_p^N D_{p,n}^N(1)} \right) \end{aligned} \quad (7.7)$$

$$- \sum_{p=k}^{n-1} \left(\frac{\eta_p^N D_{p,n}^N(\varphi_n)}{\eta_p^N D_{p,n}^N(1)} - \frac{\eta_{p+1}^N D_{p+1,n}^N(\varphi_n)}{\eta_{p+1}^N D_{p+1,n}^N(1)} \right) \left(\frac{\eta_{p+1}^N D_{p+1,n}^N(t_k)}{\eta_{p+1}^N D_{p+1,n}^N(1)} - \frac{\eta_p^N D_{p,n}^N(t_k)}{\eta_p^N D_{p,n}^N(1)} \right) \quad (7.8)$$

For the first equality, note that $\eta_n^N D_{n,n}^N(dx_{0:n}) = \mathbb{Q}_n^N(dx_{0:n})$. It is straightforward to establish that

$$\eta_p^N D_{p,n}^N(dx_{0:n}) / \eta_p^N (g(y_p | \cdot)) = \tilde{\eta}_{p+1}^N D_{p+1,n}^N(dx_{0:n}), \quad (7.9)$$

which is due to

$$\begin{aligned}
& \frac{\eta_p^N(dx_p)}{\eta_p^N(g(y_p|\cdot))} \prod_{j=p}^{n-1} Q_{j+1}(x_j, dx_{j+1}) \\
&= \frac{\eta_p^N(dx_p)g(y_p|x_p)f(x_{p+1}|x_p)}{\eta_p^N(g(y_p|\cdot)f(x_{p+1}|\cdot))} \frac{dx_{p+1}\eta_p^N(g(y_p|\cdot)f(x_{p+1}|\cdot))}{\eta_p^N(g(y_p|\cdot))} \prod_{j=p+1}^{n-1} Q_{j+1}(x_j, dx_{j+1}) \\
&= M_{p+1}^N(x_{p+1}, dx_p) \tilde{\eta}_{p+1}^N(dx_{p+1}) \prod_{j=p+1}^{n-1} Q_{j+1}(x_j, dx_{j+1}).
\end{aligned}$$

Thus

$$\begin{aligned}
& \frac{\eta_p^N D_{p,n}^N(dx_{0:p+1}, dx_n)}{\eta_p^N D_{p,n}^N(1)} - \frac{\eta_{p+1}^N D_{p+1,n}^N(dx_{0:p+1}, dx_n)}{\eta_{p+1}^N D_{p+1,n}^N(1)} \\
&= \frac{\tilde{\eta}_{p+1}^N D_{p+1,n}^N(dx_{0:p+1}, dx_n)}{\tilde{\eta}_{p+1}^N D_{p+1,n}^N(1)} - \frac{\eta_{p+1}^N D_{p+1,n}^N(dx_{0:p+1}, dx_n)}{\eta_{p+1}^N D_{p+1,n}^N(1)} \\
&= \left(\frac{\tilde{\eta}_{p+1}^N(dx_{p+1})Q_{p+1,n}(1)(x_{p+1})}{\tilde{\eta}_{p+1}^N Q_{p+1,n}(1)} - \frac{\eta_{p+1}^N(dx_{p+1})Q_{p+1,n}(1)(x_{p+1})}{\eta_{p+1}^N Q_{p+1,n}(1)} \right) \mathcal{M}_{p+1}^N(x_{p+1}, dx_{0:p}) \frac{Q_{p+1,n}(x_{p+1}, dx_n)}{Q_{p+1,n}(1)(x_{p+1})}. \tag{7.10}
\end{aligned}$$

In the first line, variables $x_{p+2:n-1}$ of the measures $\eta_p D_{p,n}^N(dx_{0:n})$ and $\eta_{p+1} D_{p+1,n}^N(dx_{0:n})$ are integrated out while the second line follows from (7.9). Using (7.10), the term (7.7) can be expressed as

$$\begin{aligned}
& \sum_{p=k}^{n-1} \int \left(\frac{\tilde{\eta}_{p+1}^N(dx_{p+1})Q_{p+1,n}(1)(x_{p+1})}{\tilde{\eta}_{p+1}^N Q_{p+1,n}(1)} - \frac{\eta_{p+1}^N(dx_{p+1})Q_{p+1,n}(1)(x_{p+1})}{\eta_{p+1}^N Q_{p+1,n}(1)} \right) \\
& \times P_{p+1,n} \left(\varphi_n - \frac{\eta_p^N D_{p,n}^N(\varphi_n)}{\eta_p^N D_{p,n}^N(1)} \right) (x_{p+1}) \mathcal{M}_{p+1}^N \left(t_k - \frac{\tilde{\eta}_{p+1}^N D_{p+1,n}^N(t_k)}{\tilde{\eta}_{p+1}^N D_{p+1,n}^N(1)} \right) (x_{p+1})
\end{aligned}$$

Note that by (3.3), (7.2) and (7.3),

$$\begin{aligned}
& \left| P_{p+1,n} \left(\varphi_n - \frac{\eta_p^N D_{p,n}^N(\varphi_n)}{\eta_p^N D_{p,n}^N(1)} \right) (x_{p+1}) \right| \leq \beta \left(\frac{Q_{p+1,n}(x_{p+1}, dx_n)}{Q_{p+1,n}(1)(x_{p+1})} \right), \\
& \left| \mathcal{M}_{p+1}^N \left(t_k - \frac{\tilde{\eta}_{p+1}^N D_{p+1,n}^N(t_k)}{\tilde{\eta}_{p+1}^N D_{p+1,n}^N(1)} \right) (x_{p+1}) \right| \leq C\beta (M_{p+1}^N \dots M_{k+1}^N).
\end{aligned}$$

Thus by (7.2) and Lemma 7.6, we conclude that there exists a finite constant C_r (depending only on r)

$$\begin{aligned}
& \sum_{p=k}^{n-1} \sqrt{N} \mathbb{E} \left\{ \left| \int \left(t_k(x_{k-1}, x_k) - \frac{\eta_p^N D_{p,n}^N(t_k)}{\eta_p^N D_{p,n}^N(1)} \right) \left(\varphi_n(x_n) - \frac{\eta_p^N D_{p,n}^N(\varphi_n)}{\eta_p^N D_{p,n}^N(1)} \right) \right. \right. \\
& \quad \left. \left. \times \left(\frac{\eta_p^N D_{p,n}^N(dx_{0:n})}{\eta_p^N D_{p,n}^N(1)} - \frac{\eta_{p+1}^N D_{p+1,n}^N(dx_{0:n})}{\eta_{p+1}^N D_{p+1,n}^N(1)} \right) \right|^r \right\}^{\frac{1}{r}} \leq (n-k) C_r \bar{\rho}^{n-k} \tag{7.11}
\end{aligned}$$

For the term (7.8), it follows from (7.10)

$$\begin{aligned} & \frac{\eta_{p+1}^N D_{p+1,n}^N(t_k)}{\eta_{p+1}^N D_{p+1,n}^N(1)} - \frac{\eta_p^N D_{p,n}^N(t_k)}{\eta_p^N D_{p,n}^N(1)} \\ &= \int \frac{\eta_{p+1}^N(dx_{p+1}) Q_{p+1,n}(1)(x_{p+1})}{\eta_{p+1}^N Q_{p+1,n}(1)} \left(\mathcal{M}_{p+1}^N(t_k)(x_{p+1}) - \frac{\tilde{\eta}_{p+1}^N(Q_{p+1,n}(1) \mathcal{M}_{p+1}^N(t_k))}{\tilde{\eta}_{p+1}^N Q_{p+1,n}(1)} \right). \end{aligned}$$

Thus, using (3.3) and (7.3), there exists some non-random constant C such that the following bound holds almost surely for all integers $k \leq p < n$, N :

$$\left| \frac{\eta_{p+1}^N D_{p+1,n}^N(t_k)}{\eta_{p+1}^N D_{p+1,n}^N(1)} - \frac{\eta_p^N D_{p,n}^N(t_k)}{\eta_p^N D_{p,n}^N(1)} \right| \leq C \bar{\rho}^{p-k+1}.$$

Combine this bound with Lemma 7.3 to conclude that there exists a finite (non-random) constant C_r (depending only on r) such that for all integers $k \leq p < n$, N :

$$\sqrt{N} \mathbb{E} \left\{ \left| \left(\frac{\eta_p^N D_{p,n}^N(\varphi_n)}{\eta_p^N D_{p,n}^N(1)} - \frac{\eta_{p+1}^N D_{p+1,n}^N(\varphi_n)}{\eta_{p+1}^N D_{p+1,n}^N(1)} \right) \left(\frac{\eta_{p+1}^N D_{p+1,n}^N(t_k)}{\eta_{p+1}^N D_{p+1,n}^N(1)} - \frac{\eta_p^N D_{p,n}^N(t_k)}{\eta_p^N D_{p,n}^N(1)} \right) \right|^r \right\}^{\frac{1}{r}} \leq C_r \bar{\rho}^{n-k} \quad (7.12)$$

The result now follows from (7.11) and (7.12).

Lemma 7.9 Assume (A). For any $r \geq 1$ there exists a constant C_r such that for all θ , y , $0 \leq k < n$, $N \geq 1$, $\varphi_n \in \text{Osc}_1(\mathcal{X})$,

$$\begin{aligned} & \sqrt{N} \mathbb{E}_\theta^y \left\{ \left| \int \frac{\eta_{\theta,k}^N D_{\theta,k,n}^N(dx_{0:n})}{\eta_{\theta,k}^N D_{\theta,k,n}^N(1)} t_{\theta,k}(x_{k-1}, x_k) \left(\varphi_n(x_n) - \frac{\eta_{\theta,k}^N D_{\theta,k,n}^N(\varphi_n)}{\eta_{\theta,k}^N D_{\theta,k,n}^N(1)} \right) \right. \right. \\ & \quad \left. \left. - \int \mathbb{Q}_{\theta,n}(dx_{0:n}) t_{\theta,k}(x_{k-1}, x_k) (\varphi_n(x_n) - \eta_{\theta,n}(\varphi_n)) \right|^r \right\}^{\frac{1}{r}} \leq C_r \bar{\rho}^{n-k} \quad (7.13) \end{aligned}$$

Proof:

$$\begin{aligned} & \int \frac{\eta_k^N D_{k,n}^N(dx_{0:n})}{\eta_k^N D_{k,n}^N(1)} t_k(x_{k-1}, x_k) \left(\varphi_n(x_n) - \frac{\eta_k^N D_{k,n}^N(\varphi_n)}{\eta_k^N D_{k,n}^N(1)} \right) \\ &= \int \mathbb{Q}_n(dx_{0:n}) t_k(x_{k-1}, x_k) (\varphi_n(x_n) - \eta_n(\varphi_n)) \\ & \quad + \int \left(\frac{\eta_k^N D_{k,n}^N(dx_{0:n})}{\eta_k^N D_{k,n}^N(1)} - \mathbb{Q}_n(dx_{0:n}) \right) t_k(x_{k-1}, x_k) (\varphi_n(x_n) - \eta_n(\varphi_n)) \quad (7.14) \end{aligned}$$

$$+ \left(\eta_n(\varphi_n) - \frac{\eta_k^N D_{k,n}^N(\varphi_n)}{\eta_k^N D_{k,n}^N(1)} \right) \frac{\eta_k^N D_{k,n}^N(t_k)}{\eta_k^N D_{k,n}^N(1)} \quad (7.15)$$

To study the errors, term (7.14) may be decomposed as

$$\begin{aligned} & \int \left(\frac{\eta_k^N D_{k,n}^N(dx_{0:n})}{\eta_k^N D_{k,n}^N(1)} - \mathbb{Q}_n(dx_{0:n}) \right) t_k(x_{k-1}, x_k) (\varphi_n(x_n) - \eta_n(\varphi_n)) \\ &= \sum_{p=0}^k \int \left(\frac{\eta_p^N D_{p,n}^N(dx_{0:n})}{\eta_p^N D_{p,n}^N(1)} - \frac{\tilde{\eta}_p^N D_{p,n}^N(dx_{0:n})}{\tilde{\eta}_p^N D_{p,n}^N(1)} \right) t_k(x_{k-1}, x_k) (\varphi_n(x_n) - \eta_n(\varphi_n)) \end{aligned}$$

with the convention that $\tilde{\eta}_0^N = \Phi_0(\eta_{-1}^N) = \eta_0$. The term corresponding to $p = k$ can be expressed as

$$\int \left(\frac{\eta_k^N(dx_k)Q_{k,n}(1)(x_k)}{\eta_k^N Q_{k,n}(1)} - \frac{\tilde{\eta}_k^N(dx_k)Q_{k,n}(1)(x_k)}{\tilde{\eta}_k^N Q_{k,n}(1)} \right) M_k^N(x_k, dx_{k-1}) t_k(x_{k-1}, x_k) P_{k,n}(\varphi_n - \eta_n(\varphi_n))(x_k)$$

Using Lemma 7.1 and Remark 7.2,

$$\begin{aligned} & \sqrt{N} \mathbb{E} \left\{ \left| \int \left(\frac{\eta_k^N(dx_k)Q_{k,n}(1)(x_k)}{\eta_k^N Q_{k,n}(1)} - \frac{\tilde{\eta}_k^N(dx_k)Q_{k,n}(1)(x_k)}{\tilde{\eta}_k^N Q_{k,n}(1)} \right) M_k^N(t_k)(x_k) P_{k,n}(\varphi_n - \eta_n(\varphi_n))(x_k) \right|^r \right\}^{\frac{1}{r}} \\ & \leq C_r \bar{\rho}^{n-k} \end{aligned}$$

Similarly, the p th term when $p < k$ can be expressed as

$$\begin{aligned} & \int \left(\frac{\eta_p^N D_{p,n}^N(dx_{0:n})}{\eta_p^N D_{p,n}^N(1)} - \frac{\tilde{\eta}_p^N D_{p,n}^N(dx_{0:n})}{\tilde{\eta}_p^N D_{p,n}^N(1)} \right) t_k(x_{k-1}, x_k) (\varphi_n(x_n) - \eta_n(\varphi_n)) \\ & = \int \left(\frac{\Phi_{p,k-1}(\eta_p^N)(dx_{k-1})Q_{k-1,n}(1)(x_{k-1})}{\Phi_{p,k-1}(\eta_p^N)Q_{k-1,n}(1)} - \frac{\Phi_{p,k-1}(\tilde{\eta}_p^N)(dx_{k-1})Q_{k-1,n}(1)(x_{k-1})}{\Phi_{p,k-1}(\tilde{\eta}_p^N)Q_{k-1,n}(1)} \right) \\ & \quad \times \int \frac{Q_k(x_{k-1}, dx_k)Q_{k,n}(1)(x_k)}{Q_{k-1,n}(1)(x_{k-1})} t_k(x_{k-1}, x_k) P_{k,n}(\varphi_n - \eta_n(\varphi_n))(x_k) \end{aligned}$$

Using Lemma 7.6 for the outer integral (recall $\Phi_{p,k-1}(\tilde{\eta}_p^N) = \Phi_{p-1,k-1}(\eta_{p-1}^N)$),

$$\begin{aligned} & \sqrt{N} \mathbb{E} \left\{ \left| \int \left(\frac{\eta_p^N D_{p,n}^N(dx_{0:n})}{\eta_p^N D_{p,n}^N(1)} - \frac{\Phi_p(\eta_{p-1}^N) D_{p,n}^N(dx_{0:n})}{\Phi_p(\eta_{p-1}^N) D_{p,n}^N(1)} \right) t_k(x_{k-1}, x_k) (\varphi_n(x_n) - \eta_n(\varphi_n)) \right|^r \right\}^{\frac{1}{r}} \\ & \leq C_r \bar{\rho}^{n-k} \bar{\rho}^{k-1-p} \end{aligned}$$

Combining both cases for p yields

$$\begin{aligned} & \sqrt{N} \mathbb{E} \left\{ \left| \int \left(\frac{\eta_k^N D_{k,n}^N(dx_{0:n})}{\eta_k^N D_{k,n}^N(1)} - \mathbb{Q}_n(dx_{0:n}) \right) t_k(x_{k-1}, x_k) (\varphi_n(x_n) - \eta_n(\varphi_n)) \right|^r \right\}^{\frac{1}{r}} \\ & \leq C_r \bar{\rho}^{n-k} \sum_{p=0}^{k-1} \bar{\rho}^{k-1-p} + C_r \bar{\rho}^{n-k} \\ & \leq C_r \bar{\rho}^{n-k} \left(1 + \frac{1}{1 - \bar{\rho}} \right). \end{aligned} \tag{7.16}$$

For (7.15), Lemma 7.5 yields the following estimate

$$\sqrt{N} \mathbb{E} \left\{ \left| \left(\eta_n(\varphi_n) - \frac{\eta_k^N D_{k,n}^N(\varphi_n)}{\eta_k^N D_{k,n}^N(1)} \right) \frac{\eta_k^N D_{k,n}^N(t_k)}{\eta_k^N D_{k,n}^N(1)} \right|^r \right\}^{\frac{1}{r}} \leq C_r \bar{\rho}^{n-k}. \tag{7.17}$$

The proof is completed by summing the bounds in (7.16), (7.17) and inflating constant C_r appropriately.

7.1 Proof of Theorem 3.1

$$\begin{aligned}\zeta_n^N(\varphi_n) - \zeta_n(\varphi_n) &= \sum_{k=0}^n \int \mathbb{Q}_n^N(dx_{0:n}) t_k(x_{k-1}, x_k) (\varphi_n(x_n) - \eta_n^N(\varphi_n)) \\ &\quad - \int \mathbb{Q}_n(dx_{0:n}) t_k(x_{k-1}, x_k) (\varphi_n(x_n) - \eta_n(\varphi_n)).\end{aligned}$$

To prove the theorem, it will be shown that the error due to the k -th term in this expression is

$$\sqrt{N} \mathbb{E} \left\{ \left| \int \mathbb{Q}_n^N(dx_{0:n}) t_k(x_{k-1}, x_k) (\varphi_n(x_n) - \eta_n^N(\varphi_n)) - \int \mathbb{Q}_n(dx_{0:n}) t_k(x_{k-1}, x_k) (\varphi_n(x_n) - \eta_n(\varphi_n)) \right|^r \right\}^{\frac{1}{r}} \leq (n-k+1) C_r \bar{\rho}^{n-k}$$

where constant C_r depends only on r and the bounds in Assumption (A) (through the estimates $\bar{\rho}$ and $\rho^2 \delta^2$ in (7.2) as well as the bounds on the score).

$$\begin{aligned}& \int \mathbb{Q}_n^N(dx_{0:n}) t_k(x_{k-1}, x_k) (\varphi_n(x_n) - \eta_n^N(\varphi_n)) - \int \mathbb{Q}_n(dx_{0:n}) t_k(x_{k-1}, x_k) (\varphi_n(x_n) - \eta_n(\varphi_n)) \\ &= \int \mathbb{Q}_n^N(dx_{0:n}) t_k(x_{k-1}, x_k) (\varphi_n(x_n) - \eta_n^N(\varphi_n)) \\ &\quad - \int \frac{\eta_k^N D_{k,n}^N(dx_{0:n})}{\eta_k^N D_{k,n}^N(1)} t_k(x_{k-1}, x_k) \left(\varphi_n(x_n) - \frac{\eta_k^N D_{k,n}^N(\varphi_n)}{\eta_k^N D_{k,n}^N(1)} \right)\end{aligned}\tag{7.18}$$

$$\begin{aligned}& + \int \frac{\eta_k^N D_{k,n}^N(dx_{0:n})}{\eta_k^N D_{k,n}^N(1)} t_k(x_{k-1}, x_k) \left(\varphi_n(x_n) - \frac{\eta_k^N D_{k,n}^N(\varphi_n)}{\eta_k^N D_{k,n}^N(1)} \right) \\ &\quad - \int \mathbb{Q}_n(dx_{0:n}) t_k(x_{k-1}, x_k) (\varphi_n(x_n) - \eta_n(\varphi_n))\end{aligned}\tag{7.19}$$

The proof is completed by summing the bounds in Lemma 7.8 for (7.18) and Lemma 7.9 for (7.19) and inflating constant C_r appropriately.

7.2 Proof of Theorem 3.2

The following result which characterizes the asymptotic behavior of the local sampling errors defined in (3.1) is proved in Del Moral [2004, Theorem 9.3.1]

Lemma 7.10 *Let $\{\varphi_n\}_{n \geq 0} \subset \mathcal{B}(\mathcal{X})$. For any $\theta, y, n \geq 0$, the random vector $(V_{\theta,0}^N(\varphi_0), \dots, V_{\theta,n}^N(\varphi_n))$ converges in law, as $N \rightarrow \infty$, to $(V_{\theta,0}(\varphi_0), \dots, V_{\theta,n}(\varphi_n))$ where $V_{\theta,i}$ is defined in (3.4).*

The following multivariate fluctuation theorem first proved under slightly different assumptions in Del Moral et al. [2010] is needed. See also Douc et al. [2009] for a related study.

Theorem 7.11 *Assume (A). For any $\theta, y, n \geq 0$, $F_n \in \mathcal{B}(\mathcal{X}^{n+1})$, $\sqrt{N} \left(\mathbb{Q}_{\theta,n}^N - \mathbb{Q}_{\theta,n} \right) (F_n)$ converges in law, as $N \rightarrow \infty$, to the centered Gaussian random variable*

$$\sum_{p=0}^n V_{\theta,p} \left(G_{\theta,p,n} \frac{D_{\theta,p,n}(F_n - \mathbb{Q}_{\theta,n}(F_n))}{D_{\theta,p,n}(1)} \right).$$

where $V_{\theta,p}$ is defined in (3.4).

Proof:

Let

$$\gamma_n = \prod_{k=0}^{n-1} \eta_k(g(y_k|\cdot))$$

and define the unnormalized measure

$$\Gamma_n = \gamma_n \mathbb{Q}_n.$$

The corresponding particle approximation is $\Gamma_n^N = \gamma_n^N \mathbb{Q}_n^N$ where $\gamma_n^N = \prod_{k=0}^{n-1} \eta_k^N(g(y_k|\cdot))$. The result is proven by studying the limit of $\sqrt{N}(\Gamma_n^N - \Gamma_n)$ since

$$[\mathbb{Q}_n^N - \mathbb{Q}_n](F_n) = \frac{1}{\gamma_n^N} [\Gamma_n^N - \Gamma_n](F_n - \mathbb{Q}_n(F_n)).$$

Note that Lemma 7.4 implies γ_n^N converges almost surely to γ_n . The key to studying the limit of $\sqrt{N}(\Gamma_n^N - \Gamma_n)$ is the decomposition

$$\sqrt{N}[\Gamma_n^N - \Gamma_n](F_n) = \sum_{p=0}^n \gamma_p^N V_p^N(D_{p,n}(F_n)) + R_n^N(F_n)$$

where the remainder term is

$$R_n^N(F_n) := \sum_{p=0}^n \gamma_p^N V_p^N(F_{p,n}^N) \quad \text{and the function} \quad F_{p,n}^N := [D_{p,n}^N - D_{p,n}](F_n)$$

By Slutsky's lemma and by the continuous mapping theorem (see van der Vaart [1998]) it suffices to show that $R_n^N(F_n)$ converges to 0, in probability, as $N \rightarrow \infty$. To prove this, it will be established that $\mathbb{E}(R_n^N(F_n)^2)$ is $\mathcal{O}(N^{-1})$. Since

$$\mathbb{E}\{R_n^N(F_n)^2\} = \sum_{p=0}^n \mathbb{E}\left\{\left(\gamma_p^N V_p^N(F_{p,n}^N)\right)^2\right\},$$

and $|\gamma_p^N| \leq c_p$ almost surely, where c_p is some non-random constant which can be derived using (A), it suffices to prove that $\mathbb{E}(V_p^N(F_{p,n}^N)^2)$ is $\mathcal{O}(N^{-1})$. By expanding the square one arrives at

$$\mathbb{E}\left(V_p^N(F_{p,n}^N)^2 \mid \mathcal{F}_{p-1}^N\right) \leq \Phi_p(\eta_{p-1}^N)\left((F_{p,n}^N)^2\right).$$

By Assumption (A), for any $x_{p-1} \in \mathcal{X}$,

$$\Phi_p(\eta_{p-1}^N)\left((F_{p,n}^N)^2\right) \leq \rho^2 \int dx_p f(x_p|x_{p-1}) F_{p,n}^N(x_p)^2.$$

By Lemma 7.7, $\mathbb{E}(V_p^N(F_{p,n}^N)^2)$ is $\mathcal{O}(N^{-1})$.

The next lemma is needed to quantify the variance of the particle estimate of the filter gradient computed using the path-based method. Note that this lemma does not require the hidden chain to be mixing. We refer the reader to Del Moral and Miclo [2001] for a propagation of chaos analysis.

For any θ , $y = \{y_n\}_{n \geq 0}$, let $\{\mathcal{V}_{\theta,n}\}_{n \geq 0}$ be a sequence of independent centered Gaussian random fields defined as follows. For any sequence of functions $\{F_n \in \mathcal{B}(\mathcal{X}^{n+1})\}_{n \geq 0}$ and any $p \geq 0$, $\{\mathcal{V}_{\theta,n}(F_n)\}_{n=0}^p$ is a collection of independent zero-mean Gaussian random variables with variances given by

$$\mathbb{E}_{\theta}(F_n(X_{0:n})^2 | y_{0:n-1}) - \mathbb{E}_{\theta}(F_n(X_{0:n}) | y_{0:n-1})^2. \quad (7.20)$$

Lemma 7.12 Let $\{\delta_\theta\}_{\theta \in \Theta} \subset [1, \infty)$ and assume $\delta_\theta^{-1} \leq g_\theta(y|x) \leq \delta_\theta$ for all $(x, y, \theta) \in \mathcal{X} \times \mathcal{Y} \times \Theta$. For any $\theta, y, n \geq 0, F_n \in \mathcal{B}(\mathcal{X}^{n+1}), \sqrt{N} (p_\theta^N(dx_{0:n}|y_{0:n-1}) - \mathbb{Q}_{\theta,n})(F_n)$ converges in law, as $N \rightarrow \infty$, to the centered Gaussian random variable

$$\sum_{p=0}^n \mathcal{V}_{\theta,p} (G_{\theta,p,n} F_{\theta,p,n}).$$

where $G_{\theta,p,n}$ was defined in (1.7) and

$$F_{\theta,p,n} = \mathbb{E}_\theta(F(X_{0:n})|x_{0:p}, y_{p+1:n-1}) - \mathbb{Q}_{\theta,n}(F_n)$$

7.2.1 Proof of Theorem 3.2

It follows from Algorithm 1 that

$$\begin{aligned} & (\zeta_n^N - \zeta_n)(\varphi_n) \\ &= \mathbb{Q}_n^N(\varphi_n T_n) - \mathbb{Q}_n(\varphi_n T_n) + \mathbb{Q}_n(\varphi_n) \mathbb{Q}_n(T_n) - \mathbb{Q}_n^N(\varphi_n) \mathbb{Q}_n^N(T_n) \end{aligned} \quad (7.21)$$

The second term on the right hand side of the equality can be expressed as

$$\begin{aligned} & \mathbb{Q}_n(\varphi_n) \mathbb{Q}_n(T_n) - \mathbb{Q}_n^N(\varphi_n) \mathbb{Q}_n^N(T_n) \\ &= \mathbb{Q}_n(\varphi_n \mathbb{Q}_n(T_n) + \mathbb{Q}_n(\varphi_n) T_n) - \mathbb{Q}_n^N(\varphi_n \mathbb{Q}_n(T_n) + \mathbb{Q}_n(\varphi_n) T_n) \\ &+ (\mathbb{Q}_n^N(\varphi_n) - \mathbb{Q}_n(\varphi_n)) (\mathbb{Q}_n(T_n) - \mathbb{Q}_n^N(T_n)). \end{aligned} \quad (7.22)$$

Combining the two expressions in (7.21) and (7.22) gives

$$\begin{aligned} & (\zeta_n^N - \zeta_n)(\varphi_n) \\ &= \mathbb{Q}_n^N((\varphi_n - \mathbb{Q}_n(\varphi_n))(T_n - \mathbb{Q}_n(T_n))) \\ &- \mathbb{Q}_n((\varphi_n - \mathbb{Q}_n(\varphi_n))(T_n - \mathbb{Q}_n(T_n))) \\ &+ (\mathbb{Q}_n^N(\varphi_n) - \mathbb{Q}_n(\varphi_n)) (\mathbb{Q}_n(T_n) - \mathbb{Q}_n^N(T_n)) \end{aligned}$$

Using Lemma 7.4 with $r = 2$ and Chebyshev's inequality, we see that $(\mathbb{Q}_n^N(\varphi_n) - \mathbb{Q}_n(\varphi_n))$ converges in probability to 0. Theorem 7.11 can now be invoked with Slutsky's theorem to arrive at the stated result in (3.5).

Moving on to the uniform bound on the variance, let

$$\begin{aligned} T_n - \mathbb{Q}_n(T_n) &= \sum_{k=0}^n \tilde{t}_k, \\ \tilde{t}_k &= t_k - \mathbb{Q}_n(t_k), \\ \tilde{\varphi}_n &= \varphi_n - \mathbb{Q}_n(\varphi_n). \end{aligned}$$

Also, the argument of V_p can be expressed as

$$\phi_p(x_p) = \frac{Q_{p,n}(1)(x_p)}{\eta_p Q_{p,n}(1)} \sum_{k=0}^n \frac{D_{p,n}(\tilde{\varphi}_n \tilde{t}_k - \mathbb{Q}_n(\tilde{\varphi}_n \tilde{t}_k))(x_p)}{D_{p,n}(1)(x_p)}.$$

It is straightforward to see that $\eta_p(\phi_p) = 0$. Therefore the variance (see (3.4)) now simplifies to

$$\text{var} \sum_{p=0}^n V_p \left(G_{p,n} \frac{D_{p,n}(F_n - \mathbb{Q}_n(F_n))}{D_{p,n}(1)} \right) = \sum_{p=0}^n \eta_p(\phi_p^2). \quad (7.23)$$

Consider the function ϕ_p . For $p \leq k-1$,

$$\begin{aligned} \frac{D_{p,n}(\tilde{\varphi}_n \tilde{t}_k - \mathbb{Q}_n(\tilde{\varphi}_n \tilde{t}_k))(x_p)}{D_{p,n}(1)(x_p)} &= \int \frac{\eta_p(dx'_p) Q_{p,n}(1)(x'_p)}{\eta_p Q_{p,n}(1)} \\ &\times \int \left(\frac{Q_{p,k-1}(x_p, dx_{k-1}) Q_{k-1,n}(1)(x_{k-1})}{Q_{p,n}(1)(x_p)} \right. \\ &\quad \left. - \frac{Q_{p,k-1}(x'_p, dx_{k-1}) Q_{k-1,n}(1)(x_{k-1})}{Q_{p,n}(1)(x'_p)} \right) \\ &\times \int \frac{Q_k(x_{k-1}, dx_k) Q_{k,n}(1)(x_k)}{Q_{k-1,n}(1)(x_{k-1})} \tilde{t}_k(x_{k-1}, x_k) P_{k,n}(\tilde{\varphi}_n)(x_k). \end{aligned}$$

Using the estimates in (3.3) and (7.2), this function is bounded by

$$\sup_{x_p} \left| \frac{D_{p,n}(\tilde{\varphi}_n \tilde{t}_k - \mathbb{Q}_n(\tilde{\varphi}_n \tilde{t}_k))(x_p)}{D_{p,n}(1)(x_p)} \right| \leq C \bar{\rho}^{n-1-p} \quad (7.24)$$

for some constant C . When $p \geq k$,

$$\begin{aligned} &\frac{D_{p,n}(\tilde{\varphi}_n \tilde{t}_k - \mathbb{Q}_n(\tilde{\varphi}_n \tilde{t}_k))(x_p)}{D_{p,n}(1)(x_p)} \\ &= \int \frac{\eta_p(dx'_p) Q_{p,n}(1)(x'_p)}{\eta_p Q_{p,n}(1)} (\mathcal{M}_p(\tilde{t}_k)(x_p) P_{p,n}(\tilde{\varphi}_n)(x_p) - \mathcal{M}_p(\tilde{t}_k)(x'_p) P_{p,n}(\tilde{\varphi}_n)(x'_p)). \end{aligned}$$

Again using the estimates in (3.3), (7.2) and (7.3),

$$\sup_{x_p} \left| \frac{D_{p,n}(\tilde{\varphi}_n \tilde{t}_k - \mathbb{Q}_n(\tilde{\varphi}_n \tilde{t}_k))(x_p)}{D_{p,n}(1)(x_p)} \right| \leq C \bar{\rho}^{n-k}. \quad (7.25)$$

Combining (7.24) and (7.25),

$$\sup_{x_p} |\phi_p(x_p)| \leq \frac{C \bar{\rho}^{n-p}}{1 - \bar{\rho}} + C \bar{\rho}^{n-p-1}(n-p),$$

$0 \leq p \leq n$. Combining this bound with (7.23) will establish the result.

References

- A. Beskos, D. Crisan, and A. Jasra. On the stability of sequential monte carlo methods in high dimensions. 2011. URL <http://arxiv.org/abs/1103.3965v1>.
- O. Cappé, E. Moulines, and T. Rydén. *Inference in Hidden Markov Models*. New York: Springer-Verlag, 2005.
- F. Cérou, F. Le Gland, and N. J. Newton. Stochastic particle methods for linear tangent filtering equations. In *Optimal Control and PDE's - Innovations and Applications*, pages 231–240. IOS Press, Amsterdam, 2001.
- P. Coquelin, R. Deguest, and R. Munos. Particle filter based policy gradient in pomdps. In *NIPS*, pages 337–344, 2008.

- P. Del Moral. *Feynman-Kac Formulae: Genealogical and Interacting Particle Systems with Applications*. New York: Springer-Verlag, 2004.
- P. Del Moral and A. Doucet. On a class of genealogical and interacting Metropolis models. In *Lecture Notes in Mathematics* **1832**, pages 415–46. Berlin: Springer-Verlag, 2003.
- P. Del Moral and L. Miclo. Branching and interacting particle systems approximations of feynman-kac formulae with applications to non-linear filtering. In J. Azma, M. Ledoux, M. Émery, and M. Yor, editors, *Séminaire de Probabilités XXXIV*, volume 1729 of *Lecture Notes in Mathematics*, pages 1–145. Springer Berlin / Heidelberg, 2000. URL <http://dx.doi.org/10.1007/BFb0103798>.
- P. Del Moral and L. Miclo. Genealogies and increasing propagations of chaos for feynman-kac and genetic models. *Annals of Applied Probability*, 11:1166–1198, 2001.
- P. Del Moral and E. Rio. Concentration inequalities for mean field particle models. Technical Report HAL-INRIA RR-6901, INRIA, 2009. URL <http://hal.inria.fr/inria-00375134/fr/>.
- P. Del Moral, A. Doucet, and S. S. Singh. A backward particle interpretation of Feynman-Kac formulae. *Ann. Inst. Stat. Math. ESAIM - Mathematical Modelling and Numerical Analysis*, 44: 947–975, 2010.
- R. Douc, A. Garivier, E. Moulines, and J. Olsson. On the forward filtering backward smoothing particle approximations of the smoothing distribution in general state spaces models. *Ann. Appl. Probab. (to appear)*, 2009. URL <http://adsabs.harvard.edu/abs/2009arXiv0904.0316D>.
- A. Doucet and V. B. Tadić. Parameter estimation in general state-space models using particle methods. *Ann. Inst. Stat. Math.*, 55:409–422, 2003.
- A. Doucet, N. de Freitas, and N. Gordon, editors. *Sequential Monte Carlo Methods in Practice*. New York: Springer-Verlag, 2001.
- J. Durbin and S. J. Koopman. *Time Series Analysis by State-Space Methods*. Cambridge University Press, 2001.
- N. J. Gordon, D. Salmond, and A. F. M. Smith. Novel approach to nonlinear/non-Gaussian Bayesian state estimation. *IEE Proc. Radar Sig. Proces.*, 140:107–113, 1993.
- K. Heine and D. Crisan. Uniform approximations of discrete-time filters. *Adv. in Appl. Probab.*, 40 (4):979–1001, 2008.
- N. Kantas, A. Doucet, S. S. Singh, and J. M. Maciejowski. An overview of sequential Monte Carlo methods for parameter estimation on general state space models. In *15th IFAC Symposium on System Identification*, volume 15, 2009. doi: 10.3182/20090706-3-FR-2004.0297.
- F. Le Gland and M. Mevel. Recursive estimation in hidden Markov models. In *36th IEEE Conf. on Decision and Control*, pages 3468–3473, 1997. doi: 10.1109/CDC.1997.652384.
- J. Olsson, R. Douc, O. Cappé, and É. Moulines. Sequential Monte Carlo smoothing with application to parameter estimation in nonlinear state space models. *Bernoulli*, 14:155–79, 2008.
- N. Oudjane and S. Rubenthaler. Stability and uniform particle approximation of nonlinear filters in case of non ergodic signals. *Stochastic Analysis and Applications*, 23:421–448, 2005.
- G. Pflug. *Optimization of stochastic models: the interface between simulation and optimization*. Boston: Kluwer Academic Publishers, 1996.

- M. K. Pitt and N. Shephard. Filtering via simulation: auxiliary particle filter. *J. Am. Statist. Ass.*, 94:590–9, 1999.
- G. Poyiadjis, A. Doucet, and S.S. Singh. Particle methods for optimal filter derivative: application to parameter estimation. In *Acoustics, Speech, and Signal Processing, 2005. Proceedings. (ICASSP '05). IEEE International Conference on*, volume 5, pages v/925 – v/928 Vol. 5, March 2005. doi: 10.1109/ICASSP.2005.1416456.
- G. Poyiadjis, A. Doucet, and S.S. Singh. Particle approximations of the score and observed information matrix in state-space models with application to parameter estimation. Technical Report CUED/F-INFENG/TR628, Cambridge University Engineering Department, 2009. URL <http://publications.eng.cam.ac.uk/16157/>.
- G. Poyiadjis, A. Doucet, and S.S. Singh. Particle approximations of the score and observed information matrix in state-space models with application to parameter estimation. *Biometrika.*, 98: 65–80, 2011.
- D. M. Titterton. Recursive parameter estimation using incomplete data. *J. Royal Statist. Soc. B*, 46:257–267, 1984.
- A. W. van der Vaart. *Asymptotic Statistics*. New York: Cambridge University Press, 1998.
- N. Whiteley. Sequential monte carlo samplers: error bounds and insensitivity to initial conditions. 2011. URL <http://arxiv.org/abs/1103.3970>.